*ASSESSING PREFERENCE FOR REINFORCERS USING
DEMAND CURVES, WORK-RATE FUNCTIONS, AND EXPANSION PATHS*

R. DON TUSTIN

A behavioral economic model that explains the choice and allocation of work rate is used to predict performance patterns in three contexts: with single schedules, with concurrent schedules when total reinforcement is low, and with concurrent schedules when reinforcement increases. Performance in the three contexts is predicted to change in orderly ways depending on how the subject evaluates the reinforcers earned. Quadrant diagrams are used to generate reinforcer demand functions, work-rate supply functions, and reinforcement-rate expansion paths. Preference between reinforcers is viewed as being a variable, with preference reversing in some situations.

Key words: preference, constrained optimization, reinforcer evaluation function, work-rate evaluation, isovalue function, demand curve, work-rate supply, expansion path, revealed preference

Behavioral economic analyses have been used to understand how schedules of reinforcement influence operant performance. One advance involves the concept of a reinforcer demand curve, which relates rate of obtained reinforcement to schedule requirement or price when a single schedule is available (Hursh, 1984). This paper extends previous analyses by predicting performance patterns when different reinforcers are available in three schedule arrangements: single schedules, concurrent schedules when total reinforcement is low, and concurrent schedules as reinforcement increases. Although previous analyses have explained performance in some of these situations, no previous analysis has predicted performance under all three schedule arrangements. Testable predictions are made using isovalue functions, which illustrate how a subject's evaluations of reinforcers can be made observable. Patterns of performance are predicted to vary in orderly ways according to the subject's evaluation of a reinforcer. The principle of revealed preference is used to deduce a subject's evaluations of reinforcers from consistent patterns of performance.

The model assumes that schedule performance involves an exchange, as responding or work is exchanged for reinforcers. It is assumed that subjects evaluate the events that are exchanged by assigning positive values to some events and negative values to other

events. It is further assumed that subjects are able to rank the relative values of events so as to form preferences that are stable when conditions are invariant. Preference is assumed to be a variable that is influenced by other variables, rather than being a result of static properties of reinforcers. Because preference is a variable, suitable measures of preference are required. Three measures of preference are proposed for different schedule arrangements.

BEHAVIORAL ECONOMIC APPROACHES

Behavioral economic approaches explain performance using constrained optimization models that assume that subjects attempt to maximize the value of certain variables (called *choice variables*) while acting under restrictions imposed by other variables (called *constraints*). Choice variables are controlled by the subject, and may be said to reflect the subject's motivation. Constraints are limitations that either are introduced by the experimenter or are inherent to a subject.

Behavioral economic models vary according to the choice variables and constraints they assume. Rachlin, Kagel, and Battalio (1980) and Rachlin, Battalio, Kagel, and Green (1981) focused on how subjects allocate time between schedules, producing time-allocation models. On the other hand, Staddon and Motheral (1978), Battalio, Green, and Kagel (1981), and Tustin and Morgan (1985) offered models that explain the allocation of work between schedules; thus, work rate was the choice variable.

Each of these models assumes that the

The contribution of Peter Morgan in developing the ideas presented in this paper is acknowledged with gratitude. Any errors remain the responsibility of the author.

Correspondence can be addressed to Don Tustin, P.O. Box 427 Glenelg, Adelaide 5045, Australia.

schedule of reinforcement is a constraint, because the schedule determines the relationship between the subject's performance and the rate of reinforcement that is delivered. In work-rate models, the subject emits work and receives reinforcers in return. The schedule sets the rate of exchange between work and reinforcement, which may be expressed in either of two ways. A schedule can be seen as defining a wage rate, in that it sets the rate of reinforcement (R) delivered for each unit of work rate (W) (wage = R/W). Alternatively, a schedule defines the price of a reinforcer, in that it sets the number of responses that must be emitted to gain a unit of the reinforcer (price = W/R). Price is the inverse of wage.

The present paper extends one work-rate model that assumes that a subject's total work rate is itself a variable (Tustin & Morgan, 1985) by examining choices involving qualitatively different reinforcers. Other models have assumed either that total work rate is constant (Rachlin & Burkhard, 1978; Staddon & Motheral, 1978) or that subjects exchange reinforcers for leisure (Battalio, Green, & Kagel, 1981). Only variable work-rate models predict both the choice of a total work rate and the allocation of work between alternative schedules (Tustin & Morgan, 1985). Work rate (W) is defined as the product of rate of instrumental responding (I) and work (w) required to emit each response (work rate = Iw), where work is defined in turn as the product of force (F) and distance (d) ($w = Fd$).

The next section illustrates how the constrained optimization approach is applied in a work-rate model.

A Work-Rate Model

The choice variables assumed in the work-rate model are work rate and reinforcement rate. It is assumed that a subject will maximize the overall value of those combinations of work and reinforcement that are available given the schedule constraints. This paper discusses applications when three simplifying assumptions are made: (a) All work is negatively valued, (b) consequences are reinforcers that are positively valued, and (c) schedules require an exchange of units of work in return for units of reinforcers. Although the model is still ap-

plicable when these assumptions are varied, space prevents discussion of these variations.

The relation set by a schedule between rate of work and rate of reinforcement is plotted in a graph whose axes are rate of reinforcement (R) and rate of work (W). Schedules are objective variables that can be drawn in the graph. Fixed-ratio (FR) schedules can be represented accurately by straight lines. However, functions for variable-interval (VI) schedules are considerably more complex (Morgan & Tustin, 1992).

ISOVALUE FUNCTIONS

Before a subject's choice variables can be represented graphically, it is necessary to transform the subject's evaluations of different combinations of work and reinforcement from subjective evaluations into objective variables. This transformation is achieved using the concept of an isovalue function. An isovalue function is a set of combinations of choice variables that are given the same overall evaluation by a subject (Green, Kagel, & Battalio, 1987; Thurstone, 1931). This paper uses two different isovalue functions. The first set of isovalue functions involves combinations of work and reinforcement that are given the same value by a subject; these are called work-reinforcement or W - R isovalue functions. The second set of isovalue functions involves combinations of two reinforcers that are given the same value by the subject; these are called reinforcer-reinforcer or R - R isovalue functions.

The question arises of how one may identify the shape of an isovalue function. The shape of a subject's isovalue function for combinations of work and reinforcement can be derived from the shapes of hypothetical functions that describe the subject's evaluations of increasing rates of reinforcement and increasing rates of work, respectively. This derivation has not previously been reported.

Reinforcer Evaluation Functions

A function that describes how a subject evaluates increasing rates of a reinforcer is called a reinforcer evaluation function. Figure 1 illustrates several possible shapes of reinforcer evaluation functions. The upper panel of Figure 1 shows functions for Reinforcers 1 and 2. The function for Reinforcer 1 is

Fig. 1. Reinforcer evaluation functions are shown for five reinforcers. Objective rate of reinforcement (R) is plotted against the subject's hypothesized evaluations of these rates of reinforcement (V_R).

higher than the function for Reinforcer 2, showing that units of Reinforcer 1 are valued more highly than units of Reinforcer 2. The two functions have similar shapes, in that they increase rapidly and then level off when moving to the right, giving the function a gentle curvature. A gently curved shape is interpreted as showing that the first few units of each reinforcer are highly valued, with increasing units being given a positive value, but a positive value that is less than the value of the earlier units. In other words, a negatively accelerated curve reflects a pattern of diminishing marginal evaluation of increasing units of the reinforcer.

Two further reinforcer evaluation functions are shown in the middle panel of Figure

1 for Reinforcers 3 and 4. The functions for Reinforcers 3 and 4 increase very steeply when the first few units of the reinforcers are provided, showing that the initial units of these reinforcers are very highly valued. However, turning points are reached, at which point the functions level off and become almost flat. A flat section in a reinforcer evaluation function shows that additional units of the reinforcer are given an almost zero evaluation by the subject.

The lower panel of Figure 1 shows yet another function, this time for Reinforcer 5. The function for Reinforcer 5 increases steadily in an almost linear manner, with each additional unit of the reinforcer being given a similar positive value to previously obtained units. In other words, Reinforcer 5 shows very little effect of diminishing marginal evaluation, because the evaluation of units of the reinforcer is almost independent of the rate at which the reinforcer is obtained.

Figure 1 illustrates shapes for reinforcer evaluation functions that differ in three ways: (a) the curvature or acceleration of evaluation functions, (b) the location of any turning point, and (c) asymptotic levels of functions. Reinforcers are said to be qualitatively different if their evaluation functions under the same levels of deprivation differ on any of these indices.

Characteristics of a reinforcer evaluation function can be interpreted. Because the curvature of a reinforcer evaluation function measures the marginal increase in value to the subject when a further unit of a reinforcer is obtained, the function assesses whether evaluations of marginal units of a reinforcer are affected by the existing level of the reinforcer. If all units of a reinforcer are given a similar evaluation regardless of the current rate of reinforcement, then the evaluation function for the reinforcer will be linear. On the other hand, if a subject's evaluation of an additional unit of reinforcement depends on the rate at which the reinforcer is delivered, then the evaluation function will be curved, showing an effect of diminishing evaluations of increasing rates of reinforcement.

A straight evaluation function means that there is no satiation effect as the rate of a reinforcer increases. On the other hand, a turning point in a reinforcer evaluation function indicates that a marked change in eval-

Fig. 2. A subject's work-rate evaluation function (W_1) for Response 1. Work rate (W) is plotted against the subject's hypothesized evaluations of these work rates (V_w).

uation of units of the reinforcer occurs at this point. The steep section of the function shows that lower rates of the reinforcer are very highly valued by the subject. A turning point may occur because a certain rate of the reinforcer is required to satisfy a "need" for the reinforcer. The needed rate of reinforcement is given by the location of the turning point. Once the needed rate is achieved, higher rates of the reinforcer are given a low evaluation, producing the plateau.

Asymptotic levels of functions show the evaluations given to very high rates of a reinforcer. Reinforcers may reach an asymptote at different levels. In particular, a reinforcer with a sharply turning evaluation function (showing that low rates of the reinforcer are highly valued) may reach an asymptote at a lower level than another reinforcer that has no satiation effect.

The shape of a reinforcer evaluation function may be affected by motivational variables. Any intervention that increases the value of a reinforcer to a subject will raise the reinforcer evaluation function. For example, if a subject is deprived of water, then the evaluation function for water will be higher and will have both a higher asymptote and a sharper turning point than if the subject were sated with water.

Work-Rate Evaluation Functions

Figure 2 illustrates one possible shape of a work-rate evaluation function (W_1) for a given response. The work-rate evaluation function falls below the origin, showing that all units of work are negatively valued as is implied by the law of least effort (Solomon, 1948). The

function has a gentle downward slope, showing that higher work rates are given progressively more negative evaluations. Although only one work-rate function is used to illustrate the model, the model applies equally if differently shaped work-rate functions are assumed.

W-R Isovalue Functions

Figure 3 shows how reinforcer evaluation functions and work-rate evaluation functions are combined to produce a work-reinforcer ($W-R$) isovalue function. Figure 3 uses a quadrant diagram to combine a subject's evaluation functions for work and reinforcement, a technique that has not previously been reported.

Figures 1 and 2 showed the subject's evaluations of choice variables on one axis. The task now is to remove these unobservable variables from the analysis. Panel A of Figure 3 reproduces a reinforcer evaluation function, and Panel B reproduces a work-rate evaluation function. Panel C uses a four-quadrant diagram to show how these two evaluation functions can be transformed so as to yield an isovalue function. The axes in Figure 3C are the reinforcement rate (R), work rate (W), and the subject's evaluations of reinforcement rates (V_R) and work rates (V_w).

The aim of the exercise is to generate a function in R/W space that reflects the shapes of the functions in R/V_R space and W/V_w space. One begins in the upper-right segment of Panel C, in V_R/V_w space, which represents the subject's evaluations of the two choice variables. A straight line ($V = V'$) is drawn in V_R/V_w space, passing through the origin. The straight line indicates that all points on the line are equally valued, so this is an isovalue line in V_R/V_w space. Because the isovalue line passes through the origin, it is valued at $V = 0$ units. The task is to find combinations of work rate and reinforcement rate that the subject also values at 0 units. The task is accomplished by moving both horizontally and vertically away from any point on the isovalue line in V_R/V_w space into the two adjacent quadrants until one intersects the evaluation function, and then moving into the lower left quadrant representing R/W space until the lines intersect.

Two example points are illustrated, labeled 1 and 2. If the process is repeated several

Fig. 3. A procedure for transforming the shape of a subject's isovalue curve from valuation space (V_R/V_W) to reinforcement-rate/work-rate (R/W) space is illustrated. Panel A shows objective rates of reinforcement (R) plotted against a subject's evaluations of these rates of reinforcement (V_R). Panel B shows objective rates of work (W) plotted against the subject's evaluations of these work rates (V_W). Panel C is a four-quadrant diagram in which the axes are V_R , V_W , R , and W . The dotted lines show how points are transformed from an isovalue curve $V = V'$ in V_R/V_W space to an isovalue function in R/W space. Panel D reproduces the isovalue curve in R/W space.

times, one establishes a set of work rates and reinforcement rates that is given the same value of 0 units, which reflects the subject's evaluation functions for the particular response and reinforcer. In Panel C, points are transposed from the upper right quadrant into the lower left quadrant where the axes are reinforcement rate (R) and work rate (W). Because both reinforcement rate and work rate are observable quantities, the exercise has identified the shape of a subject's isovalue function in terms of observable variables.

Panel D redraws the isovalue function that was derived in Panel C in a more convenient way, with reinforcement rate on the vertical axis and work rate on the horizontal axis. The isovalue function $V = V'$ has a positive slope with a positive acceleration as W increases.

The shape of a W - R isovalue function is interpreted as follows. A positive slope in a W - R isovalue function shows that the subject maintains a constant overall value by exchanging units of work for units of reinforcement. This is compatible with the view that reinforcers are positively valued and work is negatively valued. The slope of a W - R isovalue function at any point gives the rate of exchange of work for reinforcers that maintains a constant value for the subject. Curvature in a W - R isovalue function shows that the equal-value rate of exchange of work for reinforcement is not a constant, because the equal-value exchange rate depends on current levels of work and reinforcement. The positive acceleration shows that, as work rate increases, higher rates of reinforcement are required to compensate for each additional unit increase in work rate.

Figure 3 derived the shape of a subject's W - R isovalue function from the shapes of the subject's evaluation functions for a reinforcer and a response. Rachlin and Burkhard (1978) described an alternative procedure for deriving the shape of W - R isovalue functions. Their analysis required two restrictive assumptions that are not made in the present paper: (a) that all responses have a positive value, and (b) that total output remains constant so that any increase in one response necessarily reduces another response. In the present model, the use of a quadrant diagram enabled us to identify the shapes of isovalue functions when one event was positively valued and the other event was negatively val-

Fig. 4. Panel A is a four-quadrant diagram that transforms three isovalue curves ($V_1 > V_2 > V_3$) from one space to another. The axes in Panel A are rates of reinforcement (R), rates of work (W), the subject's evaluations of reinforcement rates (V_R), and the subject's evaluations of work rates (V_W). The three isovalue curves are reproduced in Panel B, in which the axes are R and W .

ued. No restriction was assumed about total response rate.

The shapes of W - R isovalue curves differ between the analysis of Rachlin et al. (1981, Figure 7) and the present analysis, because isovalue functions have a negative slope in Rachlin's analysis and a positive slope in the present analysis. The different slopes reflect a different conceptualization of the choice problem. Rachlin et al. produced a negative slope because they assumed that subjects exchange leisure for reinforcers, where leisure is the inverse of work.

A Set of Isovale Functions

The process illustrated in Figure 3 identified the shape of a subject's isovalue function for a reinforcer and a response. Figure 4 re-

Fig. 5. Panel A is a four-quadrant diagram that identifies the shapes of iso-value curves for Reinforcers 1 and 2. Axes are rate of reinforcement (R), work rate (W), the subject's evaluations of increasing reinforcement rates (V_R), and the subject's evaluations of increasing work rates (V_W). The iso-value functions derived in Panel A are reproduced in Panel B.

peats the process, but in this case a set of W - R iso-value functions is generated that have different values, where $V_1 > V_2 > V_3$. The function V_1 is valued more highly than V_2 because it has a higher intercept in V_R/V_W space. When transformed into reinforcement-rate/work-rate space, V_1 is higher than V_2 , indicating that the subject gives a higher evaluation to combinations in which higher reinforcement rates are produced for the same work rate. Figure 4 shows how a quadrant diagram can be used to generate sets of W - R iso-value functions with different rates.

Iso-value Functions for Different Reinforcers

The next step is to extend the approach so as to identify the shapes of W - R iso-value func-

Fig. 6. Panel A is a four-quadrant diagram that identifies the shapes of iso-value curves for Reinforcers 4 and 5. See Figure 5 for further details.

tions for different reinforcers. The transformation process is repeated in Figures 5 and 6. Figure 5 shows the transformation of evaluation functions for two reinforcers, Reinforcer 1 and Reinforcer 2, where units of Reinforcer 1 are more highly valued than units of Reinforcer 2 at every rate of reinforcement. The aim is to provide iso-value functions for different reinforcers that are evaluated at the same level of $V = V_1$, both in V_R/V_W space and in R/W space. The reinforcer evaluation functions have similar shapes in R/V_R space, but the evaluation function for units of Reinforcer 1 is higher than the function for units of Reinforcer 2. When these functions are transformed into R/W space to show equally valued combinations of work and reinforcement, the iso-value function for Reinforcer 2 is higher than the iso-value function for Reinforcer 1, showing that more units are required of Reinforcer 2 than of Reinforcer 1 to provide the same value to the

subject. Note that both functions represent the same value because both functions were derived from the same isovalue function where $V = V_1$. The shapes of the two functions in R/W space are similar, except that the function for Reinforcer 1 has a lower slope. The lower slope of the function for Reinforcer 1 indicates that, at any given work rate, the subject is willing to emit a higher marginal work rate to obtain an additional unit of Reinforcer 1 than of Reinforcer 2, showing that additional units of Reinforcer 1 are more highly valued.

Note that the slope of an isovalue function in R/W space measures the number of units of the reinforcer a subject requires in return for a unit increase in work rate so as to remain at the same overall value. A lower slope indicates that fewer units of reinforcement are required to compensate for a unit increase in work rate, showing that each unit of the reinforcer is more highly valued.

Any intervention that changes the value of a reinforcer for a subject will have the effect of displacing the reinforcer evaluation function in a manner similar to that illustrated in Figure 5.

Figure 6 repeats the process for two reinforcers, Reinforcer 4 and Reinforcer 5, where the reinforcer evaluation function for Reinforcer 5 has a gentle slope and the reinforcer evaluation function for Reinforcer 4 has a sharp turning point. The two reinforcer evaluation functions cross. When transformed into R/W space, the isovalue function for Reinforcer 4 has a sharp curvature, whereas the isovalue function for Reinforcer 5 has a gentle slope. The two isovalue functions cross in R/W space.

Summary

This section explained how $W-R$ isovalue functions are derived from information about the shapes of a subject's reinforcer evaluation functions and work-rate evaluation functions. Differently shaped $W-R$ isovalue functions were generated for different reinforcers as the shapes of reinforcer evaluation functions varied. Although isovalue functions are not observable, they are plotted on a graph whose axes are observable and quantifiable. The following sections will show that observable differences in performance are predicted in differ-

ent contexts, based on the shapes of isovalue functions.

SOLUTIONS FOR SINGLE-SCHEDULE PROBLEMS

Constrained optimization approaches assume that subjects will select the point that is the most highly valued of the available choice points; this point is called the *optimal point*. The model indicates that the optimal point is determined jointly by the choice variables and by the constraints. The constraints determine which alternatives are available. The subject's evaluations of the choice variables indicate which alternatives are most highly valued.

When a single schedule is available, the subject must select one combination of work rate and reinforcement rate from all of the alternatives made available by the schedule. This section explains how a constrained optimization approach is used to solve the choice problem posed by a single schedule of reinforcement. Rachlin and Burkhard (1978) offered an explanation that is formally similar, although their model differed in important details such as relating the value of responses to their duration and measuring performance in terms of relative duration of responses rather than rate of discrete responses.

The choice variables in the model are reinforcement rate and work rate. One constraint is the schedule of reinforcement. A second constraint is the subject's maximum possible rate of work. A third constraint is the time available for responding, because this restricts total work output. The first two constraints are illustrated in Figure 7. The subject's evaluations of choice variables are represented by a set of isovalue functions equivalent to those derived in Figure 4, with $V_1 > V_2 > V_3$. The subject's maximum work capacity is shown by $\bar{W}$. The schedule represents a constraint, because the subject can choose only points on the line. Points above and below the schedule line are unattainable. The subject selects different points along the schedule line by varying response rate.

The subject's evaluation of different points is represented by the $W-R$ isovalue functions. The points on the line V_1 are valued more highly than points on the isovalue functions

Fig. 7. The solution to the choice problem posed by a single schedule is illustrated. Axes are reinforcement rate (R) and work rate (W). A fixed-ratio schedule constraint (OS) is shown together with three iso-value curves ($V_1 > V_2 > V_3$). The subject's maximum work capacity is $\bar{W}$.

V_2 and V_3 , because V_1 provides higher reinforcement rates for the same work rate. However, points on the iso-value function V_1 are unattainable, because they all lie above the schedule line. Points such as b and c on the iso-value function V_3 are attainable, but they are not the most highly valued of the available points. The subject can attain a higher value than that provided by b by increasing its response rate, thereby moving upwards along the schedule line so as to gain access to a more highly valued iso-value function. The subject is predicted to increase its overall value by moving onto more highly valued iso-value functions, until it reaches a . If the subject moves past a , then it will move onto iso-value functions with lower values.

The subject is assumed to follow a marginalist approach of evaluating the marginal gains or returns in reinforcement rate against the marginal costs of responding. If a subject is at b on the schedule line, then it will find that the positive value of the last unit of reinforcement exceeds the negative value of the last unit of work required to obtain the reinforcer. The point of equilibrium occurs at a , because this is the point at which the negative value of the last unit of work is exactly offset by the positive value of the last unit of reinforcer. At this point, the net value of the exchange is zero.

Point a is the optimal point because it is the most highly valued of the attainable points. At the optimal point, the subject obtains the reinforcement-rate/work-rate com-

bination R^*/W^* . The optimal point occurs at the point at which the schedule line is tangent to the most highly valued iso-value function.

The marginalist formulation assumes that the subject is sensitive both to the slope of the schedule function and to the slope of the iso-value function, because the optimal point is the point at which these two slopes are equal. The slope of the schedule function at any response rate measures the marginal productivity of the schedule, because it gives the marginal rate of reinforcement obtained per unit of responding. The slope of the W - R iso-value function measures the equal-value rate of exchange of work for reinforcers.

The suggestion that subjects are sensitive to marginal productivities of schedules has concerned critics for two reasons. The first issue involves concern about whether the information-processing skills of subjects enable them to discriminate marginal productivities of schedules, because animal subjects are able to learn about marginal productivities of schedules only by observing the outcomes of their own behavior (Hinson & Staddon, 1983; Prelec, 1982). The second issue involves the question of how an experimenter might determine whether a subject has attempted to select the optimal point.

It has been established that subjects working on two concurrent unequal VI schedules of reinforcement spend most of their time on the richer schedule while periodically sampling the leaner schedule (Tustin & Davison, 1979). Subjects might learn about the relative marginal productivities of schedules by comparing schedules for the numbers of responses required or times spent responding before reinforcers are delivered. Tustin and Davison examined times between changeovers on concurrent VI schedules and found that when relative reinforcement rates varied between conditions, subjects adjusted the frequency of sampling the leaner schedule while maintaining the duration of this sample constant at about 3 s. In other words, subjects adjusted to changes in the relative rate of reinforcement by varying the time spent in the richer schedule before making a changeover into the leaner schedule while responding at a constant rate on both schedules. This sampling approach enables subjects to compare concurrent schedules for numbers of rein-

forcers obtained per unit of responding or per unit of time spent responding. Comparisons are made within specific time periods or "time horizons" (Timberlake, Gawley, & Lucas, 1988).

The second issue involves procedures that will show whether subjects have identified and approximated the optimal point. To assess whether a subject has discriminated marginal productivities of schedules between experimental conditions, a fuller description of the nature of the discrimination problem presented by the schedules is required. Morgan and Tustin (1992) described one approach for assessing the difficulty of the discrimination problem posed by concurrent VI schedules. They reanalyzed data from an experiment by Hunter and Davison (1982) in which pigeons' response rates varied considerably between conditions due to changes made both to force requirements on response keys and to scheduled rates of concurrently available reinforcement. Data were obtained from 6 subjects working on five VI schedules that were presented concurrently with a second schedule. In this case, the time horizon is assumed to be equivalent to an experimental session.

Morgan and Tustin (1992) plotted all rates of reinforcement obtained from each VI schedule against the response rates that earned the reinforcers. Lines of best fit were computed, producing functions that described relations between responding and reinforcement from each schedule. The marginal productivities of schedules for any given response rate are given by the slopes of these schedule functions. Although Morgan and Tustin were able to compute lines of best fit, they noted that a number of reinforcement rates fell a considerable distance away from the line of best fit for most response rates (see their Figure 1). Because reinforcement rates varied considerably for any given response rate, it is more difficult for subjects to estimate the marginal productivity of the schedule by assessing the rate of reinforcement obtained from any given response rate. The question arises as to whether subjects are able to solve these discrimination problems.

Morgan and Tustin (1992) used the concept of efficiency to assess the dual questions of (a) whether subjects were able to solve the discrimination problems and (b) whether subjects allocated responding so as to maximize

the overall reinforcement rate obtained from their chosen work rate. Performance is said to be efficient if subjects achieve the highest possible reinforcement rate from their chosen work rate. Efficient performance can be achieved only if subjects first discriminate the marginal productivities of schedules and then allocate responses so as to maximize reinforcers obtained per unit of work. Morgan and Tustin used the theoretical schedule functions to compute the maximum rates of reinforcement that could be earned by the most efficient allocation across the two response keys of the total rate of work chosen by each subject. Measures of the level of efficiency achieved by each subject were made by computing differences between these theoretical maximum reinforcement rates and the reinforcement rates actually obtained. Results showed a high level of efficiency, in that 94% of allocations achieved at least 97% of the maximum available reinforcement rate. It was concluded that subjects both were able to discriminate the marginal productivities of schedules and were motivated to maximize the rate of reinforcement obtained per unit of work.

The difficulty of discriminating marginal productivities of schedules can easily be underestimated. For example, Ettinger, Reid, and Staddon (1987) commented that a VI schedule can be represented by a line that is positively accelerated with an asymptote at the scheduled maximum reinforcement rate. However, as noted by Morgan and Tustin (1992), VI schedules produce considerable variability in reinforcement rate for any given response rate, making it difficult to estimate the marginal productivities of these schedules. VI schedules are static schedules, because the properties of the schedule do not vary as a function of the subject's own behavior. Even static VI schedules provide a substantial discrimination problem for subjects. The difficulty of discriminating the marginal productivity of a schedule increases considerably if a dynamic schedule, in which the properties of the schedule are dependent on the subject's own behavior, is used. Nonetheless, some experimenters have used dynamic schedules to test the ability of subjects to discriminate the marginal productivities of schedules. For example, Ettinger et al. (1987) provided reinforcement only if subjects met

the combined requirements of two interlocking schedules.

Before one can reasonably comment on whether subjects have responded to the marginal productivities of schedules, it is necessary to examine the nature of the discrimination task required by the schedules. A number of experiments that have supposedly examined the question of whether subjects were sensitive to marginal productivities have failed to report any assessment of the marginal productivities of the schedule functions used in the experiment. To examine the nature of the discrimination problem, experimenters must display the reinforcement rates that were obtained from each response rate, so as to show both the slope of the schedule function and any variability about a line of best fit. Current findings that subjects have not adjusted their performance according to the productivities of schedule functions may indicate simply that the subjects were unable to solve the discrimination problem posed by the schedules, and may tell us nothing about the *motivation* of subjects to allocate performance according to the marginal productivities of schedules.

In summary, this section has proposed that subjects' performance can be described as a marginalist approach to select a point on a schedule line. The chosen point is predicted to be optimal, because it provides the highest possible reinforcement rate for the chosen work rate. At the optimal point, the slope of the isovalue function is equal to the slope of the schedule function. The question of whether subjects' behavior approximates the optimal point is assessed by examining concurrent schedule performance and by measuring the efficiency with which subjects obtained the maximum reinforcement rate that was available from the chosen work rate. Efficient allocations can occur only if subjects are able to discriminate the productivities of schedules.

MODELS OF REINFORCER DEMAND AND WORK SUPPLY

The previous section showed how the constrained optimization approach is used to identify the optimal point in a single-schedule situation. This section uses isovalue functions

to generate a locus of optimal points when schedule requirements change.

Figure 8A shows a set of FR schedule lines and a set of W - R isovalue curves tangent to the schedule lines, yielding a set of optimal points. Different schedule lines represent different schedule requirements ($FR\ n$), with the slope of the schedule line decreasing as the requirement increases. W - R isovalue curves have a similar shape to those derived in Figure 3.

Figure 8A shows how the loci of a subject's optimal points will change as the schedule requirement increases for a given response and reinforcer. As the schedule requirement increases from n_1 to n_4 , the optimal reinforcement rate (R^*) decreases monotonically and the optimal work rate (W^*) first increases and then decreases, as predicted by Staddon (1979) and Rachlin et al. (1981).

Figures 8B and 8C reproduce the reinforcement-rate and work-rate performance patterns from Figure 8A, plotted against schedule requirement. The schedule requirement determines both the price of a reinforcer (I/R) and the wage rate for working (R/I), where I is the rate of instrumental responding. In Figure 8B, optimal reinforcement rates are plotted against increasing schedule requirement or increasing price, giving a reinforcer demand curve. In Figure 8C, optimal work rates are plotted against increasing schedule requirement or decreasing wage rate, giving a labor or work-rate supply curve. Figure 8 illustrates the relationship between reinforcer demand curves and work-rate supply curves, as derived using isovalue curves. Other analyses have separately derived either demand curves (Hursh, 1980, 1984; Hursh & Bauman, 1987) or work-rate supply functions (Battalio, Green, & Kagel, 1981; Battalio, Kagel, & Green, 1979). The present derivation, showing the relationship between the functions, has not previously been reported.

The present analysis derives the slope of a demand curve from the intersection of isovalue functions and schedule lines. If the shapes of schedule lines are known, then it follows that the shape of an empirical demand curve reflects information about the shape of a subject's isovalue function. The principle of revealed preference can be used to deduce how a subject evaluates reinforcers,

Fig. 8. The axes of Panel A are optimal reinforcement rates (R^*) and optimal work rates (W^*). Four iso-value functions are shown tangent to schedule lines with differing requirements (n_i), generating optimal solutions for four schedules. In Panel B optimal reinforcement rates are plotted against price or schedule requirement. In Panel C optimal work rates are plotted against schedule requirement or wage rate.

based on the shape of the subject's demand curve.

The work-rate function derived in Figure 8C is bitonic, in that it first increases and then decreases as wage rate changes. A bitonic work-rate function has also been predicted by other analyses (Allison & Boulter, 1982; Battalio, Green, & Kagel, 1981; Battalio et al., 1979; Green, Kagel, & Battalio, 1982, 1987; Kagel, Battalio, Winkler, & Fisher, 1977; Rachlin et al., 1981; Rachlin & Burkhard, 1978; Staddon, 1979). A bitonic relation between response rate and reinforcement rate is not consistent with the notion that reinforcers always strengthen behavior, because the strengthening principle implies a monotonic relation between response rate and reinforcement rate (Herrnstein, 1970; McDowell & Wood, 1985). Neither is the prediction of a

bitonic relation consistent with the hypothesis that total response rate remains constant (Herrnstein, 1974, 1979). Although current data support the notion of a bitonic relation between response rate and reinforcement rate (Battalio et al., 1979; Hanson & Timberlake, 1983; Staddon, 1979; Timberlake, 1977), the relationship warrants further investigation.

The analysis is now extended to show how the shape of a subject's reinforcer evaluation function will affect the shapes of both the reinforcer demand curve and the work-rate supply curve. Demand curves and work-rate curves are derived in Figure 9 for three differently shaped reinforcer evaluation functions, when the same schedule constraints are used. Note that the price of reinforcers increases as schedule requirement increases,

Fig. 9. Solutions for choice problems involving single schedules are illustrated for Reinforcers i, ii, and iii. The first column shows differently shaped reinforcer evaluation functions for each reinforcer. Solutions to the choice problems are illustrated in the second column. Resultant reinforcer demand functions are shown in the third column, in which optimal reinforcement rates (R^*) are plotted against the prices (P) of reinforcers or schedule requirements (n). Associated work-rate supply functions are shown in the fourth column, in which optimal work rates (W^*) are plotted against wage rate or schedule requirement.

whereas wage rate decreases as schedule requirement increases.

The subject's evaluation functions for alternative reinforcers are shaped as follows: the evaluation function for Reinforcer i in Panel A has a gentle slope; the evaluation function for Reinforcer ii in Panel E is nearly rectangular with a turning point at R_i ; and the evaluation function for Reinforcer iii in Panel J is almost linear. (These differently shaped evaluation functions were introduced in Figure 1.)

Figure 9 shows how the shape of a reinforcer evaluation function affects performance patterns. The demand curve for Reinforcer i has a moderate negative slope and a moderate intercept on the R^* axis; the work-rate function is bitonic. The demand curve for Reinforcer ii is almost flat, with a

lower intercept; the work-rate function increases monotonically and peaks near the right side of the graph. The demand curve for Reinforcer iii has an intercept high on the R^* axis and a steep slope; the work-rate function increases monotonically, with a peak towards the left of the graph.

Figure 9 shows that markedly different patterns of performance are predicted for the same changes in requirement of a single schedule of reinforcement, depending on the shape of the subject's set of reinforcer evaluation functions. The predicted work-rate functions varied from a monotonically increasing function to a bitonic function and then to a monotonically decreasing function. The shapes of the associated demand curves varied greatly.

PERFORMANCE ON CONCURRENT SCHEDULES

Earlier sections have illustrated the solutions to choice problems posed by single schedules of reinforcement. This section extends the analysis to the more complex situation in which two schedules are concurrently available. When two schedules are available, the theory must explain the allocation of work or reinforcers between schedules as well as the selection of a work rate.

A subject's total response output can be controlled by procedures such as ending an experimental session when a set number of responses has been emitted or when a certain number of reinforcers has been earned. Procedures that control a subject's response output are predicted to move performance along a locus of optimal points called an *expansion path*. An expansion path can be defined either in response-rate space or in reinforcement-rate space. It has been conventional to measure preference for a reinforcer from the ratio of responses allocated to the schedule producing that reinforcer, because the ratio of reinforcers obtained is almost invariant when VI schedules are used.

One important question is whether preference remains constant as response output changes along an expansion path. Figure 10A examines this question using a quadrant diagram. The upper left quadrant shows an almost linear reinforcer evaluation function for Reinforcer 1, and the lower right quadrant shows a sharply curved reinforcer evaluation function for Reinforcer 2. Three different isovalue functions are shown, with $V''' > V'' > V'$. Isovvalue functions in $V_1 V_2$ space are straight lines with a slope of -1 , because a constant overall value is maintained by decreasing one reinforcer by the same amount that the alternative reinforcer is increased. When transformed into a graph on which the axes are the rates of Reinforcer 1 and Reinforcer 2, the more highly valued isovalue function V''' lies to the right of V'' . The isovalue curves are reproduced in reinforcement-rate space in Panel B. Note that because these isovalue functions describe choices between two reinforcers, they are called reinforcer-reinforcer ($R-R$) isovalue curves.

The slope of an $R-R$ isovalue curve measures the marginal rate of substitutability of two re-

Fig. 10. An expansion path in reinforcement-rate (R_1/R_2) space is derived in Panel A, in which Reinforcers 1 and 2 have differently shaped evaluation functions. Panel A is a four-quadrant diagram in which axes are objective rates of reinforcement (R_1 and R_2) and the subject's evaluations of these reinforcer rates (V_1 and V_2). Three isovalue lines are shown in V_1/V_2 space, with $V''' > V'' > V'$. The isovalue curves are transformed into R_1/R_2 space by following the dotted lines. The isovalue curves are reproduced in R_1/R_2 space in Panel B. An expansion path (OM) is obtained from the points of tangency between unlabeled iso-TW lines and the highest isovalue curves.

inforcers. If two reinforcers are perfect substitutes, then the isovalue function will be a straight line, because a unit of one reinforcer completely replaces a unit of the other reinforcer. $R-R$ isovalue functions will have a negative slope, because a constant value is maintained by exchanging one reinforcer for the other. In Figure 10B, the $R-R$ isovalue function is curved convexly to the origin. The convex curvature shows that as the rate of one reinforcer decreases, higher and higher incre-

ments of the other reinforcer are required to maintain a constant overall value. The degree of convexity of the curvature measures the substitutability of the reinforcers.

A set of unmarked lines is drawn tangent to the R - R isovalue curves, which represent iso-total-work-rate functions (iso-TW). An iso-TW shows how a constant number of responses may be allocated between two schedules. A fuller discussion of how expansion paths are generated using iso-TW functions is given in Tustin and Morgan (1985). In this case, iso-TWs have a slope of -1 , because every response removed from one schedule is allocated to the other schedule. Iso-TWs farther from the origin represent higher response outputs. Optimal points occur where the iso-TW lines are tangent to the highest isovalue functions, because reinforcement rates are obtained for the minimum work rate at these points. An expansion path (OM) is drawn through the points of tangency, representing the locus of optimal points that will be chosen as response output increases.

The point of interest in Figure 10 is whether the expansion path (OM) moves along a straight line radiating from the origin, as would occur if preference remains constant when response output increases. The expansion path in Figure 10 is not a straight line, because it lies closer to the R_2 axis when response output is low and then moves closer to the R_1 axis as total reinforcement rate increases. This indicates that preference will not remain constant as reinforcement rate increases in a case in which the two reinforcers have differently shaped evaluation functions. If preference is measured from the ratio of the obtained reinforcement rates, then preference is predicted to reverse as the total reinforcement rate increases; Reinforcer 2 is preferred at low reinforcement rates, and Reinforcer 1 is preferred at high reinforcement rates.

A reversal of preference is predicted for any choice situation in which the evaluation function of one reinforcer has a sharp turning point and the evaluation function of the other reinforcer is almost linear. The reinforcer with the sharply turning evaluation function is predicted to be preferred at low rates of reinforcement, but the reinforcer with a linear evaluation function is predicted to be preferred as the level of income of reinforcers increases. The reversal of prefer-

ence is a direct outcome of the differences in shapes of the evaluation functions for the two reinforcers.

The model of choice between two differing reinforcers presented here is equivalent to the economic theory of consumer demand for two goods when there is a constraint on income (Battalio, Kagel, Rachlin, & Green, 1981; Kagel, Battalio, Green, & Rachlin, 1980; Kagel, Battalio, Rachlin, & Green, 1981; Kagel et al., 1975; Rachlin, Green, Kagel, & Battalio, 1976). This paper extends previous analyses by assessing the effect on preference of changing total response output. It is predicted that preference will not always remain constant, and that preference may reverse in conditions in which reinforcers have differently shaped evaluation functions. An example of a reversal in preference is given by Tustin (1994). The prediction that preference between qualitatively different reinforcers will change and even reverse as a subject's total reinforcement rate increases is not made by other theories of performance.

CONCLUSION

This paper describes a new way to conceptualize preference. It is proposed that preference is a variable that is best measured using appropriate functions. This represents a significant change from the proposal made by Premack (1965) that preference can be accurately estimated by observing performance in a single situation in which responses are freely available.

This argument has been presented in a number of steps. Evaluation functions for reinforcers were represented graphically. Quadrant diagrams were used to combine evaluation functions for reinforcers and for work so as to produce sets of isovalue functions. Isovalue functions were then used to predict performance in a number of schedule arrangements. The choice of a work rate was predicted when a single schedule was provided. A locus of optimal points was generated as schedule requirements changed, yielding both reinforcer demand curves and work-rate supply curves. The shapes of both reinforcer demand curves and work-rate supply curves were predicted to vary as a function of the shape of the subject's isovalue function. Predictions were made about concurrent sched-

ule performance both when reinforcement was low and when reinforcement increased. It was predicted that, in some circumstances, preference between reinforcers will not be constant but will vary and may even reverse as total reinforcement increases. No other theory has predicted performance across these three schedule arrangements, and no other theory of performance has predicted a reversal of preference as reinforcement increases.

To date, these predictions remain untested, because no experiment has examined performance under the variety of conditions that are required. To test the predictions, a subject will work for two reinforcers that are expected to have differently shaped evaluation functions in three conditions: in single-schedule arrangements, in concurrent-schedule arrangements in which reinforcement is limited, and in concurrent-schedule arrangements in which total reinforcement increases progressively. The schedule requirement will vary within each condition. The type of schedule and instrumental response will remain constant. Analyses will be facilitated if simpler schedules (e.g., FR schedules) are used, so that marginal productivities can be more easily discriminated by subjects.

It is predicted that reinforcers with differently shaped evaluation functions will generate distinctive patterns of performance. Reinforcers whose evaluation functions have sharp turning points will generate demand curves with a flat slope and an intercept that is close to the free-operant rate, generate work-rate supply curves that increase as schedule requirements increase, and will be preferred in concurrent-schedule experiments when total reinforcement is low but will become less preferred as total reinforcement increases. In contrast, reinforcers whose evaluation function is almost linear are predicted to generate demand curves with a steep slope, generate work-rate functions with a steep negative slope as schedule requirements increase, and to be not preferred in concurrent-schedule experiments when reinforcement is low but to become more preferred as total reinforcement increases. A reinforcer that might be expected to have a sharp turning point is water, and a reinforcer that might be expected to have a nearly linear evaluation

function is light, for animals deprived of these reinforcers.

It was hypothesized that the principle of revealed preference can be used to deduce the shape of a subject's evaluation function from consistent patterns of performance in relevant conditions. Three measures of performance are expected to be related, because it is assumed that they are all determined by the shape of the isovalue function, which is in turn determined by the shape of the subject's evaluation function for a reinforcer. These three measures of performance are reinforcer demand curves, work-rate supply functions, and reinforcer expansion paths.

Although the present analysis is similar to previous behavioral economic analyses, there are some crucial differences. The present model uses variables that are directly observable, because subjects are assumed to exchange work for reinforcers. Some models have introduced unobservable variables, for example by assuming that subjects exchange reinforcers for leisure (Rachlin et al., 1981). Because the present model assumes that total work rate is a variable, it is able to predict changes in preference when total response output varies.

One limitation of this paper is that it does not provide any mathematical specifications to describe subjects' evaluations of reinforcers. Mathematical specifications will permit more precise predictions of the effects of different evaluations of reinforcers.

REFERENCES

- Allison, J., & Boulter, P. (1982). Wage rate, non-labor income, and labor supply in rats. *Learning and Motivation*, 13, 324-342.
- Battalio, R. C., Green, L., & Kagel, J. (1981). Income-leisure tradeoffs of animal workers. *The American Economic Review*, 71, 621-632.
- Battalio, R. C., Kagel, J. H., & Green, L. (1979). Labor supply behavior of animal workers: Towards an experimental analysis. In V. L. Smith (Ed.), *Research in experimental economics* (Vol. 1, pp. 231-253). Greenwich, CT: JAI Press.
- Battalio, R. C., Kagel, J. H., Rachlin, H., & Green, L. (1981). Commodity choice behavior with pigeons as subjects. *Journal of Political Economy*, 89, 67-91.
- Ettinger, R. H., Reid, A. H., & Staddon, J. E. R. (1987). Sensitivity to molar feedback functions: A test of molar optimality theory. *Journal of Experimental Psychology: Animal Behavior Processes*, 13, 366-375.
- Green, L., Kagel, J. H., & Battalio, R. C. (1982). Ratio schedules of reinforcement and their relation to

- economic theories of labor supply. In M. L. Commons, R. J. Herrnstein, & H. C. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts* (pp. 395-429). Cambridge, MA: Ballinger.
- Green, L., Kagel, J. H., & Battalio, R. C. (1987). Consumption-leisure tradeoffs in pigeons: Effects of changing marginal wage rates by varying amount of reinforcer. *Journal of the Experimental Analysis of Behavior*, 47, 17-28.
- Hanson, J. H., & Timberlake, W. (1983). Regulation during challenge: A general model of learned performance under schedule constraint. *Psychological Review*, 90, 261-282.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243-266.
- Herrnstein, R. J. (1974). Formal properties of the matching law. *Journal of the Experimental Analysis of Behavior*, 21, 159-164.
- Herrnstein, R. J. (1979). Derivatives of matching. *Psychological Review*, 86, 486-495.
- Hinson, J. M., & Staddon, J. E. R. (1983). Matching, maximizing, and hill-climbing. *Journal of the Experimental Analysis of Behavior*, 40, 321-331.
- Hunter, I. W., & Davison, M. C. (1982). Independence of response force and reinforcement rate on concurrent variable-interval schedule performance. *Journal of the Experimental Analysis of Behavior*, 37, 183-198.
- Hursh, S. R. (1980). Economic concepts for the analysis of behavior. *Journal of the Experimental Analysis of Behavior*, 34, 219-238.
- Hursh, S. R. (1984). Behavioral economics. *Journal of the Experimental Analysis of Behavior*, 42, 435-452.
- Hursh, S. R., & Bauman, R. A. (1987). The behavioral analysis of demand. In L. Green & J. Kagel (Eds.), *Advances in behavioral economics* (Vol. 1, pp. 117-165). Norwood, NJ: Ablex.
- Kagel, J. H., Battalio, R. C., Green, L., & Rachlin, H. (1980). Consumer demand theory applied to choice behavior of rats. In J. E. R. Staddon (Ed.), *Limits to action: The allocation of individual behavior* (pp. 237-267). New York: Academic Press.
- Kagel, J. H., Battalio, R. C., Rachlin, H., & Green, L. (1981). Demand curves for animal consumers. *The Quarterly Journal of Economics*, 46, 1-15.
- Kagel, J. H., Battalio, R. C., Rachlin, H., Green, L., Basmann, R. L., & Klemm, W. R. (1975). Experimental studies of consumer demand using laboratory animals. *Economic Inquiry*, 13, 22-38.
- Kagel, J. H., Battalio, R. C., Winkler, R. C., & Fisher, E. B., Jr. (1977). Job choice and total labor supply: An experimental analysis. *Southern Economic Journal*, 4, 13-24.
- McDowell, J. J., & Wood, H. M. (1985). Confirmation of linear system theory prediction: Rate of changes of Herrnstein's k as a function of response-force requirement. *Journal of the Experimental Analysis of Behavior*, 43, 61-71.
- Morgan, P. B., & Tustin, R. D. (1992). The perception and efficiency of labor supply choices by pigeons. *The Economic Journal*, 102, 1134-1148.
- Prelec, D. (1982). Matching, maximizing, and the hyperbolic reinforcement schedule function. *Psychological Review*, 89, 187-225.
- Premack, D. (1965). Reinforcement theory. In D. Levine (Ed.), *Nebraska symposium on motivation* (Vol. 13, pp. 123-188). Lincoln: University of Nebraska Press.
- Rachlin, H., Battalio, R., Kagel, J., & Green, L. (1981). Maximization theory in behavioral psychology. *Behavioral and Brain Sciences*, 4, 371-417.
- Rachlin, H., & Burkhard, B. (1978). The temporal triangle: Response substitution in instrumental conditioning. *Psychological Review*, 85, 22-47.
- Rachlin, H., Green, L., Kagel, J. H., & Battalio, R. C. (1976). Economic demand theory and psychological studies of choice. In G. Bower (Ed.), *The psychology of learning and motivation* (Vol. 10, pp. 129-154). New York: Academic Press.
- Rachlin, H. C., Kagel, J. H., & Battalio, R. C. (1980). Substitutability in time allocation. *Psychological Review*, 87, 355-374.
- Solomon, R. L. (1948). The influence of work on behavior. *Psychological Bulletin*, 45, 1-40.
- Staddon, J. E. R. (1979). Operant behavior as adaptation to constraint. *Journal of Experimental Psychology: General*, 108, 48-67.
- Staddon, J. E. R., & Motheral, S. (1978). On matching and maximising in operant choice experiments. *Psychological Review*, 85, 436-444.
- Thurstone, L. L. (1931). The indifference function. *Journal of Social Psychology*, 2, 139-167.
- Timberlake, W. (1977). The application of the matching law to simple ratio schedules. *Journal of the Experimental Analysis of Behavior*, 27, 215-217.
- Timberlake, W., Gawley, D. J., & Lucas, G. A. (1988). Time horizons in rats: The effect of operant control of access to future food. *Journal of the Experimental Analysis of Behavior*, 50, 405-417.
- Tustin, R. D. (1994). Preference for reinforcers under varying schedule arrangements: A behavioral economic analysis. *Journal of Applied Behavior Analysis*, 27, 597-606.
- Tustin, R. D., & Davison, M. C. (1979). Choice: Effects of changeover schedules on concurrent performance. *Journal of the Experimental Analysis of Behavior*, 32, 75-91.
- Tustin, R. D., & Morgan, P. (1985). Choice of reinforcement rates and work rates with concurrent schedules. *Journal of Economic Psychology*, 6, 109-141.

Received October 14, 1994
Final acceptance June 6, 1995