JULY 1987

LIDS-P-1680

AD-A183 224

EXPERIMENTAL INVESTIGATION OF THE BOUNDED RATIONALITY CONSTRAINT*

by

Anne-Claire Louvet
Jeff T. Casey
Alexander H. Levis

DTIC
S ELECTE D
AUG 1 4 1987

ABSTRACT

The cognitive limitation of human decisionmakers is one of the determinants of organizational performance. A basic assumption in the analytical methodology for organizational design is that bounded rationality sets an upper limit on the amount of information that can be processed. When this rate constraint is exceeded, performance should be seriously degraded. An experiment that represents a first attempt to quantify the level of cognitive workload associated with the bounded rationality constraint of humans has been designed. The results of the experiment indicate that a threshold level can be established and used in the design of multi-person experiments.

87   7 24 001

EXPERIMENTAL INVESTIGATION OF THE BOUNDED RATIONALITY CONSTRAINT *

Anne-Claire Louvet   Jeff T. Casey   Alexander H. Levis

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139

ABSTRACT

The cognitive limitation of human decisionmakers is one of the determinants of organizational performance. A basic assumption in the analytical methodology for organizational design is that bounded rationality sets an upper limit on the amount of information that can be processed. When this rate constraint is exceeded, performance should be seriously degraded. An experiment that represents a first attempt to quantify the level of cognitive workload associated with the bounded rationality constraint of humans has been designed. The results of the experiment indicate that a threshold level can be established and used in the design of multi-person experiments.

I. INTRODUCTION

Individual bounded rationality is one of the major determinants of performance of information processing and decision making organizations. This is especially true for organizations which must perform under severe time constraints (e.g., tactical military organizations). Therefore, in designing such organizations and the command, control, and communications systems which support them, individual cognitive workload[1] is of critical concern. The present study represents a first attempt to quantify the level of cognitive workload associated with the bounded rationality constraint of humans. This will be done within the context of well-structured, time-constrained decision tasks for which the decisionmaker is well-trained.

The organization of this paper is as follows. First, an information theoretic surrogate for workload is described (Boettcher and Levis, 1982). In order to use the workload surrogate in predicting organizational performance under time constraint, a critical assumption about the nature of individual bounded rationality is required. This assumption is described in terms of the Yerkes-Dodson law relating performance and arousal. Then this assumption is tested in the context of a simplified air defense task. The results are shown to provide support for the existence of a person-specific bounded rationality constraint. Individual differences in this constraint are described and prescriptive implications for organizational design are discussed.

The Workload Surrogate

Boettcher and Levis (1982) proposed a method for modeling the cognitive workload of individual decision makers. This *workload surrogate*, based on N-dimensional information theory, has three key features:

(1) it is *objective*: it does not require decision makers to form introspective judgments concerning their workloads;

(2) it is *comprehensive*: it takes into account not only the uncertainty contained in the input information, but also the uncertainty associated with plausible situation assessment and response selection algorithms;

(3) it potentially offers *inter-task comparability* within the task context specified above.[2]

The workload surrogate is computed using n-dimensional information theory. Information theory is build upon two primary quantities: entropy and transmission. Entropy is a measure of information and uncertainty: given a variable x belonging to the alphabet X, the entropy of x is :

$$H(x) = -\sum_x p(x) \log p(x) \qquad (1)$$

Entropy is measured in bits when the base of the logarithm is two.

The transmission - also called mutual information - between variables x and y, elements of X and Y, and given p(x) and p(y), and p(x|y) is defined as follows:

$$T(x:y) = H(x) - H_y(x) \qquad (2)$$

where

$$H_y(x) = -\sum_y p(y) \sum_x p(x|y) \log p(x|y) \qquad (3)$$

The expression for transmission (2) generalizes for n-dimensions to:

$$T(x_1:x_2:...x_n) = \sum_i H(x_i) - H(x_1, x_2, ... , x_n) \qquad (4)$$

The workload surrogate, denoted by G, is defined as being the total information processing activity of a system i.e., the sum of the entropies of all the variables in the system. The Partition Law of Information (PLI) (Conant, 1976) can be used to decompose the total activity G into components that characterize what may happen to information as it is processed by a system . For a system with input variable x, N-1 internal variables w_i and output variable y, the PLI states:

$$\sum_{i=1} H(w_i) = T(x:y) + T_y(x:w_1, w_2, ... w_{N-1}) \qquad (5)$$
$$+ T(w_1:w_2:...w_{N-1}) + H_x(w_1, w_2, ... w_{N-1}, y)$$

This equation may be abbreviated (keeping the same order) as:

$$G = G_t + G_b + G_c + G_n \qquad (6)$$

where G_t (called throughput) measures the amount by which the input and output are related; G_b (called blockage) is the amount of information which enters the system, but is not present in the output or blockage; G_c (or coordination) is the amount by which the internal variables interact; and G_n (or noise) is the uncertainty in the system when the input is known.

Bounded Rationality and the Yerkes-Dodson Law

The *bounded rationality constraint*, for present purposes, refers to a hypothesized characteristic of a particular region of the function relating decision making performance to cognitive workload, where workload is calculated using the workload surrogate. Considerable experimental psychological work has examined the influence of *arousal* on performance in various types of tasks (Kahneman,1973). Arousal is, in turn, influenced by a variety of factors, including cognitive workload. The commonly observed relation between performance and arousal, called the Yerkes-Dodson curve or "law", is shown in Figure 1. This relation is obtained when arousal is varied over an extremely wide range. At very low arousal, performance is low due to boredom and vigilance limitations. At very high arousal, performance is low due to extreme stress and sensory overload. However, in a well-designed organization, all decision makers should be operating near the top of the curve at all times. Thus the central region of the curve is of particular interest. It is important for an organizational designer to know how much cognitive workload (e.g., in bits of information processed per unit time) the organization members can cope with without substantial decrements in performance due to overload. If the bounded rationality constraint is to be quantified in this way, it must first be established that performance does indeed begin to drop at a predictable point with increasing workload.


Fig.1 Relationship between the Yerkes - Dodson law and T^*

Workload, Processing Rate, and Interarrival Time

Workload can be manipulated in two basic ways: by varying task complexity (amount of processing required) or by varying the amount of time available for doing a given, well defined task. The latter method was used in the present work. Consider a situation in which a decision maker is performing a

sequence of independent tasks that have interarrival time τ. Performance for each task is a binary variable, taking on the value 1, if the decision is accurate and timely, and 0 otherwise. The workload experienced by the decision maker can be expressed as:

$$G = F t'; \qquad t' \leq \tau, \qquad (7)$$

where G is the decision maker's workload per task, F is the decision maker's processing rate, and t' is the portion of the interarrival time during which the decision maker is processing the task. If τ is more than ample, various tradeoffs between F and t' are possible, any one of which will get the work done within the allotted time, τ. However, for sufficiently small values of τ, t' will approach τ, and the decision maker must increase F in order to maintain G and thereby avoid a decrease in performance.

The Threshold Hypothesis

This reasoning leads to a *threshold hypothesis* which is the focus of the paper: *As τ is decreased, a point, T^*, will be reached beyond which further increases in F are impossible. As a result, performance will drop substantially.* This hypothesis is an underlying *assumption* in the organizational design methodology proposed by Levis and his co-workers (1984). Under this assumption, the bounded rationality constraint, F_{max}, can be expressed as:

$$F_{max} = G_\tau / T^* \qquad (8)$$

where F_{max} is the upper bound on the decision maker's rate of processing in bits per unit time, G_τ is the workload (total information theoretic activity in bits) per task computed analytically via the workload surrogate, and T^* is the interarrival time threshold below which performance deteriorates substantially. T^* must be measured initially experimentally. The obtained value of T^* depends, of course, upon the structure of the task, as does the computation of G_τ.

If a T^* value can be found practically for an individual performing an experimental task, then the individual's bounded rationality constraint can be estimated quantitatively from (2). If this constraint shows reasonable stability across well defined tasks, then it is of interest to attempt to characterize individual differences in the bounded rationality constraint. Thus three key empirical issues must be addressed:

(1) Is there a person-specific threshold? That is, for most individuals, is performance uniformly high whenever the minimum (or "continuous duty") processing rate required by the task is less than some person-specific critical (F_{max}) value?

(2) Is F_{max} robust within an individual to manipulation of task parameters unrelated to workload?

(3) What are the characteristics of the distribution of F_{max} for a sample of individuals? Uni- or multi-modal? Normal or skewed? Low or high variance?

If workload per task, G_τ is assumed constant as certain task (or person) parameters are varied, then F_{max} can be constant only if T^* is found to be constant. Thus, these questions can be answered directly in terms of T^* rather than F_{max}. This simplification eliminates the need to compute G_τ

separately for each task variant or person. The assumption of constant G_r will be met whenever a single information processing algorithm (strategy) is used consistently, or the alternative algorithms degrade as a function of t in essentially the same manner. The resulting prescription for experimental design is that tasks should be constructed so as to limit possible algorithmic variability, both within and between subjects.

II. EXPERIMENTAL METHOD

The experiment involved a highly simplified tactical air defense task using an IBM PC[3]. The basic screen display is shown in Figure 2. The large circle represents a radar screen. On each of a series of trials, either four or seven incoming threats were present. Two pieces of information were provided concerning each threat: relative speed and relative distance. The information for each threat was presented as a ratio of two two-digit integers[4]. The distance was the numerator and the speed the denominator. All threats were assumed to be converging on the center of the screen.


Fig. 2 The screen display used in the experiment

The subject's task was to select the threat which would arrive first in the absence of interception. Since the ratio represents the time it will take for the threat to reach the center of the screen, the task can be interpreted as one of selecting the minimum ratio. In order to limit strategic variability, two restrictions were imposed. First, ratios were displayed in pairs; only two ratios were visible at a time. This procedure eliminated variation in the order of information acquisition. Second, to reduce the incidence of responding based on incomplete information, a final response was permitted only after all of the four or seven ratios had been displayed.

Ratios could appear only on the vertical or horizontal diameter of the radar screen. Thus ratios appeared in one of four regions: left, right, above, or below the center. Each ratio was randomly assigned to one of these four regions, subject to the constraint that no two ratios appear in the same region at the same time. For each pair of ratios, the subject indicated his or her choice by pressing one of four arrow keys corresponding to the direction of the ratio from the radar screen's center. The arrow keys were located on the numeric keypad of the keyboard and were arranged isomorphically with the four regions of the radar screen.

The physical distance of each ratio from the center was proportional to its numeric distance. However, in order to restrict strategic variability, subjects were instructed to attend only to the numeric information. This restriction is important, because Greitzer and Hershman (1984) showed that an experienced Air Intercept Controller tended to use physical distance information only (and not speed information) in determining which of a number of incoming threats to prosecute first. This simplified strategy was labelled the *range* strategy. The operator was, however, able to use both range and speed information -- the *threat* strategy -- when instructed explicitly to do so. The threat strategy, if executed in a timely way, is of course more effective than the simpler range strategy.

The ratio chosen as smallest was retained on the radar screen and the other was replaced with one of the remaining ratios. This procedure was repeated until all ratios had been examined. Row(s) of small rectangles to the left of the radar screen indicated the total number of threats for the current trial and the number yet to be examined (see Figure 2). Each time a new ratio appeared on the radar screen, one of the rectangles disappeared.

Performance feedback was provided at the end of each trial for which the subject finished the comparisons within the allotted time. In this case, only one ratio remained on the screen at the end of the trial. If this ratio was in fact the smallest, it "flashed" several times to indicate a correct response. If this ratio was not the smallest, a low-pitched tone signalled the error. This tone (which subjects reported to be particularly obnoxious) was used to reduce strategic variability by biasing subjects against guessing.

Manipulation of Interarrival Time

The amount of time allotted for each trial was shown by a fixed clock hand (see Figure 2). A moving second hand (running clockwise from 12 o'clock) indicated elapsed time within a trial. A 1.5 sec. pause prior to the start of each trial allowed subjects to see how much time was allotted. The fixed hand flashed during this interval.

Time per trial was varied in alternating descending and ascending series. Twice as much time was allotted for seven as for four threat trials, because the number of comparisons was double (six versus three). In order to retain comparability between the four and seven threat conditions, t is defined as time per comparison, rather than as trial interarrival time. t was varied from 0.75 sec. to 3.5 sec. in 0.25 sec. increments for both conditions. Thus 12 values of t were used. t ranged from 2.25 to 10.5 sec. for four threats and from 4.5 to 21 sec. for seven threats.

The task was constructed to minimize the influence on performance of time required for non-cognitive (i.e., perceptual and motor activity). Even the minimum t of 0.75 sec. allows ample time for eye movements, perception, and motor response. The limiting factor in response time then is the rate of cognitive activity, F_{max}.

Organization of Trials

The number of threats was constant within blocks of 24 trials. A block of trials consisted of a descending series over the 12 values of t, followed by an ascending series. The number of threats was then changed for the subsequent block. There was a 2.5 sec. pause between blocks, during which time the large rectangle to the left of the radar screen (see Figure 2) flashed to indicate the impending change in number of threats.

For each subject, the full experiment consisted of 24 iterations (12 descending and 12 ascending) over values of t for each number of threats. The total duration was approximately 2.5 hours. Eight iterations were completed in each of three sessions. Subjects typically participated in no more than one session per day. Prior to each experimental session, subjects were given a brief (three to five minute) "warmup" period during which no data were recorded.

Practice Session

Subjects received a 30 minute practice session prior to the actual experiment. This session consisted of six iterations over t for each number of threats. For the practice session, t was varied from 1 to 5 sec. per comparison in 0.5 sec. increments. Informal discussion with subjects indicated that most felt their performance would not improve substantially with additional practice. The practice data were not analyzed.

Subjects

Twenty-one subjects participated in the experiment. The majority of subjects were present or former MIT students. They were paid a flat rate of approximately $6 per session.

III. RESULTS AND DISCUSSION

Data Reduction and Transformation

Each subject received a binary performance score for each trial. A trial was scored one, if the three or six comparisons were completed within the allotted time and the correct threat (minimum ratio) was selected. Otherwise, the trial was scored zero. Figure 3 shows the raw data matrix of binary scores for one iteration over values of t -- that is, for one replication. The full data matrix was four-dimensional, consisting of 21 subjects by 12 values of t by 2 sets of threats by 24 replications. These binary scores were converted into proportions by summing over replications and dividing by 24. This yielded a three dimensional matrix of proportions, consisting of 21 subjects by 12 values of t by 2 sets of threats.

Proportion data violate the homogeneity of error variance assumption required for regression and curve-fitting. In order to equate these error variances, the proportion data were transformed via the formula:

$$(\sin^{-1}\sqrt{\text{proportion}})/1.57 \qquad (9)$$

where the denominator is a scaling constant. The general effect of this transformation is to increase slightly small proportions, while decreasing slightly large proportions. All analyses reported herein are based on the transformed proportion data.

The threshold hypothesis will be evaluated in terms of T* with respect to the three questions discussed above: Is there a person-specific threshold in performance as a function of the required processing rate? Is this threshold robust to minor task changes unrelated to workload (i.e., changes in number of threats with time per threat held constant)? How can the distribution of individuals' thresholds be characterized?

Is There a Person-Specific Threshold F_{max}?

General characteristics of curves: For each subject, two data curves were plotted, one for four threats and one for seven threats. Figure 4a-c shows performance as a function of interarrival time for three subjects. These curves were selected

Values of τ

Fig. 3 The data matrix for one replication

as examples of the strongest (a), typical (b), and weakest (c) degrees of support for the threshold hypothesis contained in the set of 42 (21 subjects for two sets of threats each) curves obtained from the experiment. Visual inspection of the entire set of curves revealed the following general characteristics:

(1) They do not have the Yerkes-Dodson concave shape. This indicates that the experiment succeeded in tapping into the moderate-to-high arousal portion of the Yerkes-Dodson curve (see Figure 1), rather than the "vigilance" portion.

(2) Most curves are nearly flat (zero slope) for large values of τ.

(3) They have positive slopes for smaller values of τ.

(4) Some "leveling-off" to very small positive slopes for the smallest τ values.

These characteristics suggest that some type of growth curve, also called "S" curve, would be appropriate for summarizing the data.

$$J = a\,e^{-b\,e^{-ct}} \qquad (10)$$

where J is performance expressed as a proportion.

Gompertz curve parameters were estimated independently for each of the 42 data sets. Figures 4 a-c also show the Gompertz fit superimposed on the observed data . The Gompertz fit was quite good ($0.93 \le r^2 \le 0.99$) in every case. Thus, in effect, the Gompertz curves provide a concise mathematical description of the data. Table 1 summarizes the degree of fit of the data to the S curves versus linear regression.

Table 1.   Summary of r^2 values for S vs. linear functions

	Mean	Std. dev.	Min.	Max.
S	0.98	0.01	0.93	0.99
Linear	0.75	0.15	0.25	0.94

Estimation of T^*: For each subject, two T^* values were estimated (one for each number of threats) from the appropriate S curve. As illustrated in Figure 5, T^* was taken as the t value corresponding to the intersection of two lines tangent to the S curve. One of the lines was tangent to the curve's asymptote as $t \to \infty$ while the other was tangent at the inflection point. This method for estimating T^* was chosen because it is conservative in the sense that performance is not degraded seriously, even for values of t somewhat smaller than T^*. This conservativeness results from the extrapolation to asymptotic performance (which, according to the Yerkes-Dodson law, is never reached) in defining the upper tangent line. Figure 5 shows the tangent lines and resulting T^* value for the same S curve as shown in Figure 4a.


Fig. 4 a-c   Performamce as a function of interarrival time for three subjects.

The use of a growth curve can be justified not only by plots of the experimental data, but also by the physical behaviour that underlies the growth curves. Growth curves are characterized by their S shape: the growth starts slowly (characterized by a nearly flat curve segment), then the growth increases rapidly (steep slope) and finally levels off to an optimum or saturation level (the curve flattens again). Recall that the purpose of the experiment is to investigate whether there is a bounded rationality constraint i.e., whether well trained decisionmakers under time pressure will perform near optimum until they are beyond the bounded rationality constraint when performance will decrease rapidly; as the time pressure further increases their performance will quickly fall to an almost null level. Considering the hypothesis, a growth curve seems most appropriate, since it characterizes patterns where quantities increase from near zero to the optimum level very rapidly.

For the purpose of this experiment, the most appropriate curve of the family is the Gompertz curve which has the characteristic of not being symmetric about the inflection point. This is a relevant property, since one can not predict that performance will decrease in a symmetric way after the subject is working beyond the bounded rationality constraint. Also, it was almost impossible to get a significant number of data points near null performance, since the time could not be decreased indefinitely: poor performance had to be caused by the incapacity to process mentally the task and not by physical limitations such as time needed to press the necessary keys.

The Gompertz curve has three degrees of freedom and is given by (Martino,1972):

Fig. 5 Estimate of T^* value using S curve approximation

The obtained T^* values are summarized in Table 2. The mean value of T^* over subjects and numbers of threats was 2.21 sec. (standard deviation: 0.70). Confidence in the method for estimating T^* is increased by the finding that for both threat conditions the mean value of T^* over subjects was roughly equal (i.e., within 0.05 sec.) to the T^* obtained from the curve of the mean over subjects.

Table 2. Summary of T^* values (in seconds) for 4 and 7 threats

	Mean	Std. dev.	Min.	Max.
4 threats	2.19	0.77	0.96	4.46
7 threats	2.23	0.65	1.26	3.68

Two subjects had T^* values greater than the maximum t of 3.5 sec. (up to 4.46 sec.). These subjects' curves were unique in that they had substantial positive slopes even at the maximum t. Had the maximum t been increased for these subjects, a near-zero slope region would presumably have been encountered.

The data clearly support the existence of a threshold, T^*, for each of the 21 subjects tested. This result opens the door to information theoretic quantification of the bounded rationality constraint, F_{max}, for each subject via Equation (2). The only additional information needed to solve for F_{max} would be an estimate of the workload, G_r, computed analytically using the workload surrogate.

Is the Threshold Robust to Task Changes Unrelated to Workload?

Robustness of the threshold to changes in the number of threats would help to establish that, to some degree, the bounded rationality constraint is stable across tasks. If, however, instability were found for such a minor task change, there would be no need to go further. In addition, the effect of manipulating number of threats is of some intrinsic interest, because of implications for how subjects manage their time.

Effective time management is more critical for seven than for four threats, while "overhead" or "start-up" time is more critical for four threats than for seven. Subjects knew before the start of each trial how much time was allocated for the trial. Part of the subject's task was to budget the available time over the three or six comparisons so that all comparisons could be completed and full use made of the available time. The criticality of accurate budgeting can be seen from Equation (11).

$$\text{Response Time} = n \, t' + b, \qquad (11)$$

where n is the number of comparisons (three or six), t' is the amount of time the subject allocates for each comparison, and b is the overhead, startup, or initialization time for a trial. The value of b is independent of n. According to this model, the subject must choose t' so that the resulting response time is less than or equal to n τ. Clearly, with increasing n, the detrimental effect of setting t' non-optimally increases relative to the detrimental effect of the fixed overhead, b.

Comparison of T^* for 4 threats and T^* for 7 threats: As Table 2 suggests, no systematic differences were found in T^* as a function of the number of threats. The mean over subjects for T^* (4 threats) did not differ significantly from that for T^* (7 threats), $t(19) = 0.28$, $p > 0.05$. Twelve of 21 subjects had $T^*(7) > T^*(4)$. This proportion is not significantly different from 0.5. Moreover, no subject performed significantly better for one number of threats than for the other.

The absence of systematic differences in T^* due to the number of threats manipulation, provides modest evidence that T^*, and therefore F_{max}, may be a *stable individual characteristic* within the class of well-structured, time-constrained information processing and decisionmaking tasks. This stability suggests that it may not be necessary to measure a decision maker's F_{max} value for every type of task the decision maker may have to perform. Instead, the decision maker's F_{max} value could be measured using a prototypic "calibration" task. The value obtained from this prototypic task could be safely assumed to apply to a substantial range of structurally similar tasks.

Distribution of Individual Differences in T^*

The case for the existence and stability of the bounded rationality constraint, F_{max}, within individuals is clear. T^* and, therefore, F_{max} are stable as the number of threats is varied from four to seven. In addition, these results provide indirect evidence for the stability of F_{max} over time, since each subject was tested on three or four different days. (A "composite" curve resulting from wide day to day fluctuations in the bou..ed rationality constraint would not likely reveal a clear thresh....) However, despite this intraperson stability, one would not expect interperson stability. Indeed, a finding that T^* is constant across people (even in a sample of MIT students) would be disturbingly counterintuitive. As Table 2 suggests, there is substantial variability in T^* across individuals. Figure 6 shows the frequency distribution of 21 T^* values. These values were obtained by averaging each subject's $T^*(4)$ and $T^*(7)$ value: each value is the T^* associated to one subject. This distribution is unimodal, very peaked, and has mean 2.21 sec. and standard deviation 0.70 sec. Moreover, the distribution of individual T^* values can be characterized as normal -- A χ^2 test for goodness of fit revealed non-significant deviation from normality: $Q^2 = 5.65 < \chi^2(.95,2) = 5.99$.


Fig. 6 Distribution of T^* values.

IV. CONCLUSIONS

The assumption in the organizational design methodology that an individual decision maker's performance will be adequate whenever a fixed bounded rationality constraint is not exceeded is supported by the present results. In addition, the bounded rationality constraint appears not to be affected by superficial task variables that do not appreciably change workload. Finally, across individuals, the bounded rationality constraint was found to be normally distributed.

This finding along with the finding that F_{max} is normally distributed has an important prescriptive implication for organizational design. In an organization consisting of several decisionmakers, tasks must be allocated so that no decisionmaker is overloaded. The knowledge that F_{max} is normally distributed can be used to make the first allocation of tasks; then when specific individuals are assigned to the tasks, a further fine tuning of the organization can be done by the designer. The cognitive workload and the tempo of operations can be specified so that the individual organization members operate below, but near their threshold. This leads the way for controlled experiments in which the individual workloads can be manipulated to exceed the bounded rationality constraint and the effect on the organizational performance observed. That is, with the information obtained from this experiment, it is now possible to calibrate the mathematical model of distributed tactical decisionmaking organizations and determine the range of parameters over which experiments should be carried out.

V. REFERENCES

Boettcher, K. L., and A. H. Levis (1982). Modeling the interacting desisionmaker with bounded rationality. *IEEE Trans. Sys. Man Cybern.*, SMC-12, pp. 334-344.

Conant, R. C. (1976). Laws of information which govern systems. *IEEE Trans. Sys. Man Cybern.*, SMC-6, pp. 240-255.

Greitzer, F. L., and R. L. Hershman (1984). Modeling decision making strategies in simulated AAW operations. *Proc. 7th MIT/ONR Workshop on C^3 Systems*. LIDS-R-1437, MIT, Cambridge, MA. pp. 145-150.

Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall.

Levis, A. H. (1984). Information processing and decisionmaking organizations: A mathematical description. *Large Scale Systems*. vol. 7, pp. 151-163.

Martino, J. P. (1972). *Technological Forecasting for Decision-making*. New York, NY: Elsevier

FOOTNOTES

[1] The term "workload" will be used throughout to refer specifically to cognitive workload as opposed to, for example, perceptual or manual workload.

[2] The PLI should be a very useful tool when designing multi person experiments since it considers different aspects of communication within an organization. The PLI could help identify the variables (or characteristics) of the organization which are of interest and help predict how the total workload and performance of the organization will be affected because of changes of one or more of these critical variables.

[3] The experiment was run on a Compaq Deskpro Model 2 equipped with 8087 math coprocessor, monochrome graphics card (640 X 200 pixels), 640K of memory, and monochrome monitor. Programming was in Turbo Pascal version 3.01A. The operating system was MS-DOS version 2.11. It was also run on an IBM PC AT with the 80287 math coprocessor and with 640K of memory. None of the high resolution graphics capabilities if the AT were used so that the experiment be portable to a wide variety of PC compatible machines.

[4] Speeds and distances were selected subject to the following constraints: (1) greater than 10 and less than 98, (2) no multiples of 10. Each speed and distance combination was screened and rejected if the resulting ratio was: (1) a whole number, (2) no speed value be used more than once per trial; and (3) no multiples of 10 be used. Distances were selected independently of speeds, but subject to the same constraints. For each trial, all ratios were either greater than or less than one. This restriction was included because pilot work had shown that decisions ratios on opposite sides of one were trivially easy, regardless of interarrival times. The greater-than-one / less-than-one determination was made randomly for each trial. Candidate ratios were screened against the following criteria: (1) each possible pair of ratios within a trial must differ by no less than 0.05 and by no more than 0.9, while,in the greater than one condition, the minimum allowable ratio was 1.2; (3) ratios yielding whole number quotients were not allowed. If a candidate ratio failed on any criterion, a new ratio was generated and the process repeated until a complete set of 4 or 7 compatible ratios had been obtained.

END

9-87

DTIC