

AD-A070 937

TECHNION - ISRAEL INST OF TECH HAIFA FACULTY OF INDU--ETC F/6 5/9
INTERPRETATIONS OF TASK DIFFICULTY IN TERMS OF RESOURCES: EFFIC--ETC(U)
NOV 78 D NAVON, D GOPHER

UNCLASSIFIED

AFOSR-79-0828

NL

| OF |
AD
A070937




END
DATE
FILMED
8-79
DDC

12
P.S.

LEVEL

ADA070937



המרכז לחקר התהליכים הרפואיים הפסיכופיזיולוגיים וההנדסיים
של בטיחות האדם העובד
מוסד הטכניון למחקר ופתוח בע"מ

הטכניון - מכון טכנולוגי לישראל
הפקולטה להנדסת העשיה וניהול

TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY
FACULTY OF INDUSTRIAL AND MANAGEMENT ENGINEERING

DDC
RECEIVED
JUL 9 1979
RECEIVED
A

DDC FILE COPY

HAIFA, ISRAEL

6
INTERPRETATIONS OF TASK DIFFICULTY IN TERMS OF RESOURCES:
EFFICIENCY, LOAD, DEMAND, AND COST COMPOSITION.

16 2313

17 A2

10 David Navon Daniel Gopher

15 AFOSR-78-3131

9 Technical Report AFOSR-78-1

12 46p

11 November 1978

18 AFOSR, AFOSR

19 79-4828,
78-1

Prepared for the Life Sciences Directorate
of the United States Air Force Office of Scientific Research
and European Office of Aerospace Research and Development

Technion - Israel Institute of Technology
Faculty of Industrial and Management Engineering
The Center of Human Engineering
and Industrial Safety Research

DDC
RECEIVED
JUL 9 1979
A

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (AFSC)
NOTICE OF TRANSMITTAL TO DDC
This technical report has been reviewed and is
approved for public release IAW AFR 190-12 (7b).
Distribution is unlimited.
A. D. BLOSE
Technical Information Officer
distribution unlimited.

404 498

A

JOB

ACKNOWLEDGMENT

The research presented in this report was supported by the Life Sciences Program, Air Force Office of Scientific Research, under grant no. AFOSR 78-3131. Captain Jack Thorpe in the Life Sciences Directorate and Captain Robert Powell in AFOSR European Office were the scientific monitors of this grant.

Accession For	
NTIS GRL&I	<input checked="" type="checkbox"/>
DDC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	<input type="checkbox"/>
By _____	
Distribution/	
Availability Codes	
Dist	Available/or special
A	

↓

ABSTRACT

The effect of task difficulty on performance can be conceptualized within a theory which posits that performance depends on the use of resources from a single pool. When the difficulty of a task is said to increase, it may mean either that resources invested in it can now do less (i.e. a decrease in efficiency), or are now required to do more (i.e. an increase in load), or have now less time to do it (i.e. a stricter limit on processing duration). Either way, difficulty should most often interact with resource investment in such a way that effects of resource investment on quality or speed of performance are more pronounced the easier the task is.

If the processing system is viewed as comprised of a number of mechanisms each having its own capacity, which may be considered as a separate resource, then a difficulty manipulation may affect differentially the use of each of those capacities. If in a dual-task situation a manipulation of the difficulty of one task affects the use of a mechanism which is not required by the other task, processing of the latter may remain intact under some circumstances.

To get a complete picture of how difficulty affects dual-task performance, it is proposed to manipulate task preferences as well as difficulty parameters and to present their joint effect by families of POCs. An application of this methodology to the study of pursuit tracking is briefly described and interpreted in terms of multiple resources.

A

C

Interpretations of Task Difficulty in terms of Resources:
Efficiency, Load, Demand, and Cost Composition

It often occurs that performance of a task is affected not only by its own difficulty, but also by the existence or by the difficulty of another task with which it is time-shared. This was taken to indicate that both tasks apply demands to the same capacity (or resources, effort, attention, etc. See, e.g., Broadbent, 1971; Kahneman, 1973; Kerr, 1973; Moray, 1967; Norman & Bobrow, 1975; Posner & Boies, 1971; Shiffrin, 1976) and get supplies in proportions that are related to their relative demands. The notion of capacity is widely used but there is little agreement about what it is or how to go about testing it (cf. Kantowitz, in press). In particular, if one were to extract from the literature a clear prediction from capacity models about the effect of task difficulty on concurrent task performance or about the joint effect of difficulty levels of both tasks, he would be at a loss.

In this paper we address these problems from a point of view which is based on some ideas borrowed from microeconomics, an approach which we describe and discuss in detail elsewhere (Navon & Gopher, in preparation).

As a first step let us examine how much of the confusion about predictions can be traced back to obscurity of terminology. And, in fact, the term "task difficulty" is frequently considered by different researchers to denote quite different things. A recent reminder of the disagreement about the meaning of "difficulty" is the exchange between Kantowitz & Knight and Lane. Kantowitz & Knight (1976) derived from capacity models the prediction that while performance of a difficult primary task will be

impaired by conjoining it with a secondary one (or making the secondary one more difficult), the performance of an easy primary task will show very little, if any, decrement as a result of such manipulations. They interpreted failures to find such an interaction in some studies, including their own, as an embarrassment for capacity models. Lane (1977) argued that the argument raised by Kantowitz & Knight was incorrect, because the existence of interaction of the sort they considered as a necessary prediction from capacity models depends on the shape of the function relating performance to task difficulty and available resources, and the particular choice of difficulty levels of each task. In their answer to Lane, Kantowitz & Knight (in press) contended that Lane had misinterpreted their argument because of his failure to understand properly the concept of difficulty as they used it. They argued that "difficulty" should be interpreted to denote the amount of resources required to achieve a specific level of performance of a certain task, thus, whenever a task is claimed to be more difficult than another one, it follows by definition that it actually uses more resources when performed in isolation.

We believe the source of confusion is that the word "difficulty" in its natural language usage is polysemous. To enable communication among researchers, this polysemy has to be first recognized and then resolved by defining the various senses and associating each of them with a different label. We therefore start by proposing a taxonomy.

Performance and Resources

Task difficulty can be manipulated by varying any of a number of variables such as sensory quality of stimuli, predictability of stimuli, availability and completeness of relevant memory codes, S-R compatibility,

response complexity, etc. Each one of those variables affects performance. The issue is in what manner.

To examine this issue, let us postulate that the human system possesses a limited amount of resources, such as processing facilities. The usage of those resources is the mental input the system invests to produce mental or behavioral output, such as information transmission. Since resources are always there to be used, that mental input is a lasting entity like a flow of a stream, thus may be viewed, if you will, as processing energy. However, since resources are limited, there is also a limit on the availability of that mental input per a unit time (or at a given point in time), namely on what may be viewed as processing power. It is to this processing power that we usually refer by saying that processing resources are being invested or allocated.

Processing output is positively related to the amount of mental input used to produce it, just as the number of produced shoes depends on the number of labor hours put into production (see, e.g., the output-input functions in Figure 1A). Hence, the amount of invested resources

Insert Figure 1 about here

(viz., mental input per unit time) determines output rate (viz., processing output per unit time), in much the same way that the number of workers allotted to shoe production determines the daily shoe production or that the availability of communication channels for transmitting a message determines the information transmission rate. Thus, output-input functions as in Figure 1A can be translated to rate-resource functions as in Figure 1B.




Fig.1: Illustrations for two output-input functions (panel A), their corresponding rate-resource functions (panel B), four adequacy-resource functions (panel C), and two latency-resource functions (panel D), derived from one of the rate-resource functions.

As can be seen in Figure 1A, output depends not only on how much is invested, but also on what the investment can produce. The contribution of a unit of input to the total output (or of a unit of resources to output rate) can be termed processing efficiency. If the output-input relation is not linear, then a distinction should be made between marginal efficiency which is the output gain resulting from the last unit of input (e.g., α in Figure 1A is the marginal efficiency at input I_0), and average efficiency, which is the average contribution of all units invested (e.g., β in Figure 1A is the average efficiency at I_0). It is easy to see that total output is a product of amount of invested input and average efficiency (and that output rate is a product of amount of invested resources and average efficiency).

Efficiency may vary across situations, tasks, and levels of performance. It is presumably determined by certain parameters of the task or of the subject or of their specific combination. Some of these parameters are related to general and specific abilities and skills of the subject, and some others, such as luminance-contrast or signal-to-noise ratio, are usually considered as difficulty variables (or data quality; see Norman & Bobrow, 1975). Some levels of these parameters may forestall performance completely; some others just make it more or less difficult. Thus, one sense of task difficulty is the output rate per a unit of resources (or if you prefer the reciprocal definition, the cost of a unit of output rate in terms of resources). The two curves in Figure 1A or Figure 1B correspond to two different efficiencies due to variation of luminance contrast. This argument can be rephrased in information-theoretic terms for those who prefer to think of resources as transmitting

channels. The rate of transmission is not constant; it depends on some parameters of the task; for example, a manipulation of signal-to-noise ratio affects the rate of signal transmission.

Thus far, we have not mentioned the word "performance". Would it be correct to regard either of the functions in Figures 1A or 1B as a performance-resource function in the sense used by Norman & Bobrow (1975)? Not always.

Tasks differ in the type of requirements they set for performance. In some tasks, there is no criterion in reference to which output can be assessed. This is the case, for example, with tasks which involve transmitting information with no specification of amount to be transmitted, such as talking. Performance of the latter tasks is usually defined simply in terms of output rate (say, how much information is conveyed during one minute of talking), so its dependence on invested resources is described by a rate-resource function of the type illustrated in Figure 1B.

The case is different with tasks for which the objective of processing is to attain some critical amount of output, say a correct detection response with perfect certainty. The critical level may be called task load, because it determines how much mental input should be invested to meet the criterion on output. Performance of such tasks is typically measured in one of two ways, adequacy or latency. The selection between these two greatly depends on the nature of the task.

Adequacy of performance can be defined as actual output/load, i.e. it is the ratio of actual to critical output. Since given sufficient time it is always possible to meet the criterion on output, adequacy measures are typically relevant in situations in which processing duration of a

single response is limited, either because response rate is externally paced (e.g., tapping) or because stimuli are available for processing for a brief period (e.g., recognition of tachistoscopically presented stimuli, externally-paced visual search). A continuous task such as tracking can also be classified in this category, if we construe it as a sequence of location correction responses each of which is to be executed fast enough to keep up with the change of stimuli. From the basic definition of adequacy, it is simple to derive that it can also be expressed as processing-duration \times (output rate/load), i.e. the proportion of criterial output produced at a unit time multiplied by the length of the processing period². The function relating adequacy to amount of invested resources can be derived from the underlying rate-resource function by dividing each rate by the load and multiplying it by processing duration. As an illustration consider the four functions in Figure 1C which depict the adequacy of four variants of a task whose processing efficiency is presented in Figure 1B. For example, suppose the task is to identify one of a certain type of stimuli at a high luminance contrast. If the rate of output of the identification process with R_0 resources is two units per second, and if accuracy is linearly related to amount of output, then if 4 units are required for an absolutely certain identification, percentage of correct responses with R_0 resources will be $\frac{1}{2}$ at a one second exposure-duration and $\frac{1}{4}$ at a half-second exposure duration³; if the load is 6 units, the respective percentages will be one-third and one-sixth. Note that all accuracy measures may be interpreted as adequacy values, but not only them. Every performance measure in which actual input is referred to criterial output may be considered as some transformation of adequacy (for example, proportion of hits in tapping tasks, percentage of on-target time in tracking tasks).

Sometimes processing duration is not imposed by the nature of the task (for example, in subject-paced visual search). The performer is free to continue processing until the criterial output is met. In this case the relevant aspect of performance is the latency to criterion which can be defined as $c + \text{load/output rate}$, where c is some constant period representing the contribution of factors which are unrelated to processing done with the resources in question. The function relating latency to amount of invested resources can be derived from the underlying rate-resource function if we know the values of c and task load. For example, the two functions in Figure 1D describe the latency-resource relationships for two variants of the task in Figure 1B given a high luminance contrast. When c is one second and R_0 resources are invested, then if the load is 4 units, 3 seconds will be required to reach the criterion; if the load is 6 units, 4 seconds will be required. If the performer aims at a certain constant level of output which is different from the criterial one, latency-resource functions can still be derived in a similar fashion (where load is replaced by that other level of output). But of course, it is hard to predict what will happen if the performer varies the adequacy level he aims for in an unsystematic fashion (say, shifts his position along the speed-accuracy trade-off curve).

Since both task load and processing duration affect the quality of performance produced by investment of resources (or equivalently, the cost of a unit performance) they can be rightfully considered as some aspects of task difficulty⁴. In some tasks both can be manipulated by the experimenter. For example, in an externally paced visual search, the experimenter can vary the stimulus onset asynchrony (SOA) on the one hand and the size of the memory set on the other hand⁵. This can be done

on top of efficiency manipulations which can be achieved, for example, by varying the luminance contrast.

So there are three major types of function which relate performance to amount of invested resources: rate-resource functions, adequacy-resource functions, and latency-resource functions. All of them are affected by characteristics of the task and the performer, which may be termed subject-task parameters, in the following way: all of them are affected by processing efficiency. Two are affected by task load. Just one is affected by processing duration.

When subject-task parameters are given and a certain level of performance is intended (not to be confused with the criterial level of output with respect to which performance is defined), the amount of resources required to achieve that level under the circumstances can be derived from the performance-resource function. This theoretical quantity may be called the demand for resources. Note that according to this definition of demand, demand is not an invariant property of a task as implied by the analyses of some previous authors (see, e.g., Kerr, 1973); it is rather defined for a specific task and a specific level of performance. Because performance degrades gracefully (Norman & Bobrow, 1975) task demand may be a variable quantity. It is determined not only by the objective constraints, namely by the subject-task parameters, but also by the intentions and the allocation policy of the system.

Sometimes what we call task demand is referred to as difficulty (e.g., Kantowitz & Knight, in press). This usage of the term "difficulty" is consonant with one natural language sense of this word denoting the subjective feeling of strain accompanying involvement in demanding tasks.

However, this sense of "difficulty" which refers to the total cost in resources imposed by a given level of performance of a task should be absolutely distinguished from the other sense which denotes the cost per a unit of performance improvement determined only by subject-task parameters (via efficiency, load, or processing duration). From this perspective the argument between Kantowitz & Knight (1976; in press) and Lane (1977) seems to hinge mainly on a semantic ambiguity. Each of the parties interprets "difficulty" in a different manner. The legitimacy of both interpretations cannot be denied and the issue is which one captures best the manipulations in the experimental situations to which the authors refer. We would like to skip this issue, and just note that later in this paper we use the term "task difficulty" in the cost-per-unit sense.

In sum, the effect of resource allocation on performance is determined by constraints imposed by the encounter of the specific task and the individual subject, in much the same way that the effect of labor allocation on the yield of corn is constrained by the climate, soil fertility and particular properties of the corn plant. Those constraints are the subject-task parameters. Some of them affect processing efficiency, some others affect task load, and some others a maximal processing period. All of them participate in determining the performance-resource function, where performance can be measured in terms of rate, adequacy, or latency to criterion. The system is then free to determine the task demand by intending to a certain performance level.

A Prediction

It follows from this discussion that the effect of any subject-task parameter on performance is probably multiplicative with the amount of resources. The output is the aggregate of the contributions of all units of resources invested. If an efficiency manipulation affects the marginal efficiency of all units of resources to some extent, then the effect cumulates so that the functions diverge (see Figure 1B). If the manipulation affects the marginal efficiency by the same factor (say, multiply it by two over the whole region), then the joint effect of the efficiency manipulation and amount of invested resources is multiplicative. Manipulation of load clearly affects performance multiplicatively. So does manipulation of processing duration.

Thus, it is quite safe to expect that effects of difficulty and available capacity will be multiplicative or at least interact in some similar way, namely that effects of resource investment on rate or adequacy will be more pronounced the easier the task is (and the reverse will hold for effects on latency). Note, however, that if the functions approach a ceiling dictated by the nature of the performance measure (what Norman & Bobrow, 1975, call a "data-limit"), that interaction may disappear or even be reversed in some regions of the performance functions (see Lane, 1977).

How can the amount of invested resources in a task be controlled? One common approach is to pair the task with another concurrent one (e.g., Rolfe, 1971). Another approach is to manipulate the difficulty of the concurrent one (e.g., Kantowitz & Knight, 1974; 1976). We later point to some weaknesses of both approaches and advocate another one which is

is supposed to operate via the subjects' voluntary control on their own resources. In any event, we will have to consider how a subject operates in a dual-task situation. We will briefly touch on some concepts and notions presented in detail in Navon & Gopher (in preparation).

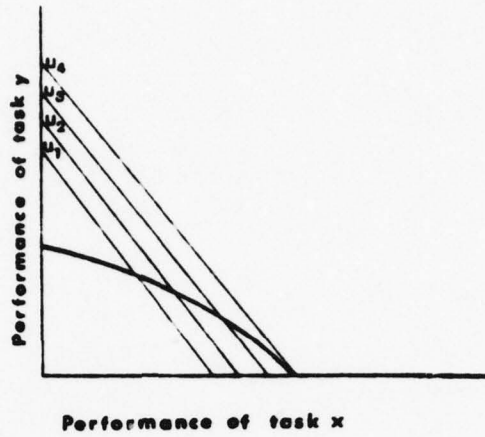
Resource Allocation

In a dual-task situation, given the structure of the tasks and the capacity of the performer, some levels of joint performance are feasible and some others are not. Most often, the performer can achieve every performance combination that can be given by the performance functions of the tasks as long as the amount of resources used by both tasks together does not exceed his capacity. The set of performance combinations that can be produced when the performer operates at his full capacity, can be represented as a curve of the type called by Norman & Bobrow (1975) performance operating characteristics (or POC in short; see bold curve in Figure 2).


Insert Figure 2 about here

POCs may have various shapes. The slope of a POC at a given point reflects the relative contribution of resources to the two tasks: a unit of resources moved from task y to task x leads to a decrease in performance of y by the marginal contribution to y and to an increase in performance of x by the marginal contribution to x.


The POC comprises of a set of alternative combinations only one of which is realized in a particular situation. So, if the performer can control the selection among the alternative combinations, he will probably



A



B



C

Fig.2: Illustrations for three types of indifference curves. Panel A describes perfect utility trade-off. Panel B describes complete lack of trade-off. Panel C describes partial substitution. Each of the two sets of curves in this panel, the solid curves and the dashed ones, correspond to a different situation or subject. The bold curve is a performance operating characteristic.

consider their utility. His preferences can be represented by means of indifference curves each of which is a locus of all combinations among which the person is indifferent. The solid and dashed thin curves in each panel of Figure 2 illustrate two maps of indifference curves. The slope of an indifference curve represents the utility trade-off of the tasks, namely how much improvement in performance of task y is needed to compensate for deterioration in performance of task x in terms of utility, that is to maintain utility at the same level. The optimal combination of performance levels is at the point of the POC which is tangent to the "north-easternmost" indifference curve (e.g., point E in Figure 2). At that point the slopes of the POC and the indifference curve are equal, which means that no extra utility can be gained by trading either more x for less y or vice-versa. If the performer aims at that optimal point when allocating his resources, then it can be concluded that resource allocation depends on both objective relative cost of tasks per a unit performance and subjective task preferences.

Thus a POC for a given pair of tasks and for a given subject may be regarded as the locus of all performance combinations which arise from splitting capacity between the two tasks in all different ratios (namely under all possible task priorities) when the performance-resource functions for both tasks are given (namely when subject-task parameters are fixed).

Since task difficulty is considered to affect the productivity of a unit of resources, namely the slope of the performance-resource function, it will also change the slope of the POC, provided that the difficulty of the concurrent task is held constant. Thus, the effect of difficulty on dual-task situations can be described by means of a family of POCs. When

task x is made more difficult, the POC has a smaller x intercept (see Figure 3). When task y is made more difficult, the y intercept is smaller.

Insert Figure 3 about here

Making both tasks more difficult should depress both intercepts.

This is how difficulty affects the set of feasible alternatives for joint performance. We now turn to discuss how it affects the actual combination selected given a certain pattern of task preferences. Does difficulty affect the demand for resources? It does, of course, affect the demand per unit performance, so that to maintain performance at the same level the performer has to respond to increased difficulty by a larger resource investment. However, it rarely happens that all aspects of performance remain intact when difficulty is increased. Thus, difficulty may be reflected in deterioration with no concomitant change in resource allocation. If adequacy (say, accuracy) is emphasized and there is no speed requirement, difficulty may affect latency. If, on the other hand, speed is emphasized, difficulty will affect adequacy. Only if both speed and adequacy are to be maintained at some desired level, the performer will have to recruit more resources to compensate for increased difficulty. So absence of difficulty effects on concurrent performance (e.g., Briggs, Peters & Fisher, 1972; Kantowitz & Knight, 1976; Wattenbarger & Pachela, 1972) is not very surprising, if it is found that the difficulty affects the performance of the task itself. Furthermore, it is often the most expected result: task difficulty is usually more likely to affect the performance of the same task rather than the performance of the concurrent one. This prediction becomes clear when one inspects the

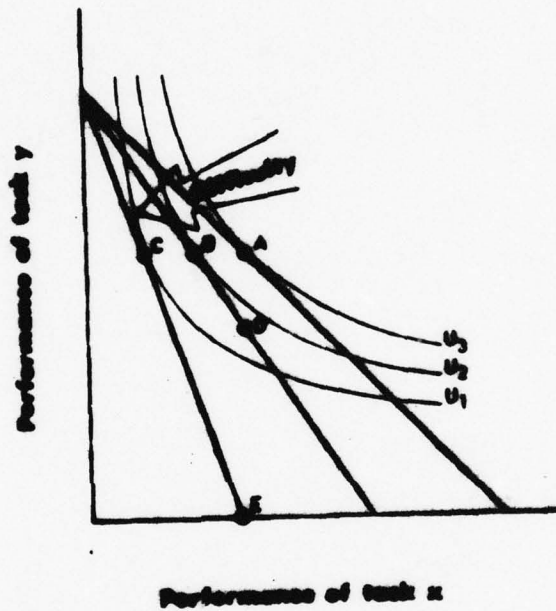


Fig.3: An illustration for the effect of task difficulty given the indifference map presented in the figure. A, B and C are points of optimal resource allocation for easy, medium, and difficult task x respectively. D and E represent joint performance under medium and high difficulty respectively when performance of task x is supposed to be protected.

illustration in Figure 3: increasing the difficulty of task x under the task preferences represented by the indifference curves shifts optimal joint performance from point A to point B, and then to point C. Of course, a precise prediction is not possible without knowledge of the exact shape of the indifference curves and the POC.

Measurement of Capacity Interference

How does one go about testing capacity interference or measuring it?

One frequent approach is to observe performance decrements from single- to dual-task situations. This is a poor indication for capacity interference, because such decrements may result from other kinds of interference which we elsewhere termed concurrency cost (see Navon & Gopher, in preparation), or perhaps may be counteracted by a tendency of capacity to stretch in order to accommodate the heavier load (see Kahneman, 1973).

Another approach is to examine the effect of the difficulty of one task on the performance of the other. The rationale is that the more difficult a task, the more it consumes resources that under the capacity interference hypothesis could otherwise have been invested in the performance of the concurrent task (Kerr, 1973). However, when applying this method one should beware of several pitfalls.

First, this rationale is valid as long as the difficulty manipulation of one task does not inadvertently affect the difficulty of the other one as well. For example, suppose a subject is required to identify a stimulus presented at a fixed location while tracking a randomly moving target to another location (see, e.g., Gopher & North, 1974). The more difficult tracking is, the higher the mean distance between the target and

the stimulus to be identified, thus if the subject fixates the target, the stimulus will be seen more peripherally.

Second, as is shown above, the effect of difficulty may be confounded with considerations of allocation policy which are hard to correct for. One remedy would be to ensure that the performance of the manipulated task is maintained at the same level. This is the logic underlying the secondary task technique, under which subjects are instructed to regard the manipulated task as primary and protect its performance against interference from the secondary one. However, such instructions may not be sufficient to attain that objective, since actual performance may nevertheless turn out to vary with difficulty (as is the case, for example, in an experiment reported by Griffith and Johnston, 1977; see also reviews by Kerr, 1973, and Rolfe, 1971). A statistical solution is to analyze the effect of difficulty on secondary task performance by means of analysis of covariance where the covariate is the corresponding primary task performance.

It should be realized that the source of the problem is the use of just one condition of resource allocation, namely a single point of a POC, for each level of difficulty. That is why the information which a procedure yields about processing potential is confounded with motivational aspects, namely with allocation policy. Thus, a more promising approach is to estimate a complete POC for every given level of difficulty, so that the difficulty effect will be manifested by a family of POCs.

How does one obtain an empirical POC? If we assume that the subject controls his own processing devices, then the experimenter can try to influence resource allocation by simply telling the subject how to do it.

In other words, experimenters should fix subject-task parameters for both tasks, allow the subjects maximal control over quality of performance for both, and induce them to change the relative emphasis on the tasks by means of pay-offs or instructions (cf. Norman & Bobrow, 1976). A family of POCs can be obtained by varying the difficulty of one task and plotting a POC for every level of difficulty.

How do we know whether subjects consume all available capacity? If they are required to perform as well as they can, and in addition we observe that task performance is related to both difficulty and processing priority, we are quite safe in assuming that subjects do not aim just at meeting a certain "satisficing" level of performance but rather try to do the best they can. Sometimes, however, not all available capacity is required to reach optimal performance (cf. the notion of "data limit" in Norman & Bobrow, 1975). Then the POC should be considered as the bound of joint performance, rather than as the locus of consequences of full capacity operation.

In sum, families of POCs serve to separate between effects of difficulty and effects of allocation policy; they are obtained by joint manipulation of task priorities and subject-task parameters.

Multiple Resources

Up to this point resources have been construed as a sort of general undifferentiated entity very much analogous to a common currency in a monetary system or to energy in a physical system or to the general intelligence factor G in theories of human intelligence; tasks interfere to the extent that they depend on resources from that general pool. Elsewhere we advanced, discussed, and reviewed some evidence for the idea that there may

be various types of resources as there are various factors that may be input to production. The human processing system may be viewed as comprised of a number of processing mechanisms each having its own capacity. Each specific capacity constrains the output rate of a specific mechanism, and it can be shared by several concurrent processes, thus it constitutes a distributable resource.

If several specific resources exist, then performance depends on the amounts of each of them. First, suppose that to perform a certain task resources are used in fixed proportions, e.g., exactly two units of STM capacity with one unit of VIS capacity; any increase in one of them without a concomitant increase in the other would not improve performance at all.

To be exact, a task is characterized by a required mixture of outputs (or output rates, in case synchronization of activity is essential) of the various resources. The proportions as well as actual amounts of the various resources needed to realize that output mixture depends on their efficiencies. Let the combination of specific resources used to obtain a unit of performance be called a cost composition⁶. Intending to a certain level of performance determines the combination of total amounts of specific resources demanded by the task which may be called a demand composition. It follows from the above definitions that the cost composition is determined both by the nature of the task and by the efficiencies of the various resources. Manipulating a subject-task parameter may theoretically affect equally the efficiencies of all relevant resources, so that the cost composition vary in terms of amounts but not in terms of proportions. In this case we may say that the manipulation affects performance quantitatively. However, we suspect that most often manipu-

lation of a subject-task parameter has a differential effect on different processing facilities, so that it changes the relative weights of the various resources in the cost composition. That may be regarded as a qualitative modification in the nature of the task.

Some types of resources are not relevant at all for certain tasks. Thus, for any task x , all the resources can be classified into two classes, the set of resources which can be used by task x (X) and the set of irrelevant resources (\bar{X}).

Different tasks may have different compositions of specific resources. Some tasks may even use resources of a type which is not used at all by other tasks. For any two tasks x and y , the whole arsenal of resources can be viewed as composed of four sets: $X \cap Y$, which is the set of resources usable by both tasks; $X - Y$, which is the set of resources that can be used by task x but not by task y ; $Y - X$, which is the set of resources that can be used by task y but not by task x ; and $\bar{X} \cap \bar{Y}$, which is the set of resources irrelevant for both tasks. Tasks interfere with each other to the extent that their cost compositions are similar so that they have to compete for some common types of resource. Elsewhere (Navon & Gopher, in preparation) we have shown that the amount of task interference reflected by the concavity of the POC is affected, amongst other things, by the existence of common resources (namely of the set $X \cap Y$), the similarity of composition of these common resources, and to what extent the pool of disjoint resources ($Y - X$ or $X - Y$) is being exhausted.

Let us consider now the implications of the notion of multiple resources for difficulty effects in dual-task situations. Suppose the performance of task y is observed to be related to the manipulation of a

certain parameter. It may mean that that manipulation affects the cost of the task in terms of some (or all) types of resources. If that task is conjoined with another one x which does not use those types of resources, the manipulation will be ineffective for the performance of the latter task. Note that this can happen in two cases. One, when $X \cap Y = \emptyset$, namely the tasks do not overlap at all in their use of resources. Two, when $X \cap Y \neq \emptyset$ but the difficulty manipulation of task y has a differential effect on the components of the cost composition, so that it taps only resources from the set $Y-X$.

The effect of the difficulty manipulation in the first case is illustrated in Figure 4A: there is no performance trade-off at all, so manipulating the difficulty of a task affects just the maximal level of performance of that task.

Insert Figure 4 about here

In the second case, however, there is some trade-off, but since the manipulation has no bearing on the cost composition of common resources, it affects neither the performance of the competing task nor the amount of trade-off. When the difficulty manipulation makes a resource of the set $Y-X$ (or $X-Y$) scarce relative to a common resource, it also restricts the amount of the common resource that can be used by the manipulated task so that the residual is made available to the other task. A simple example may illustrate this point best. Suppose the system possesses 20 units of R^1 and 20 units of R^2 . Further, suppose that the demand for resources for a unit increase of both tasks is constant over all levels of performance: task x demands one unit of R^1 ; task y demands one unit of R^1 and one unit

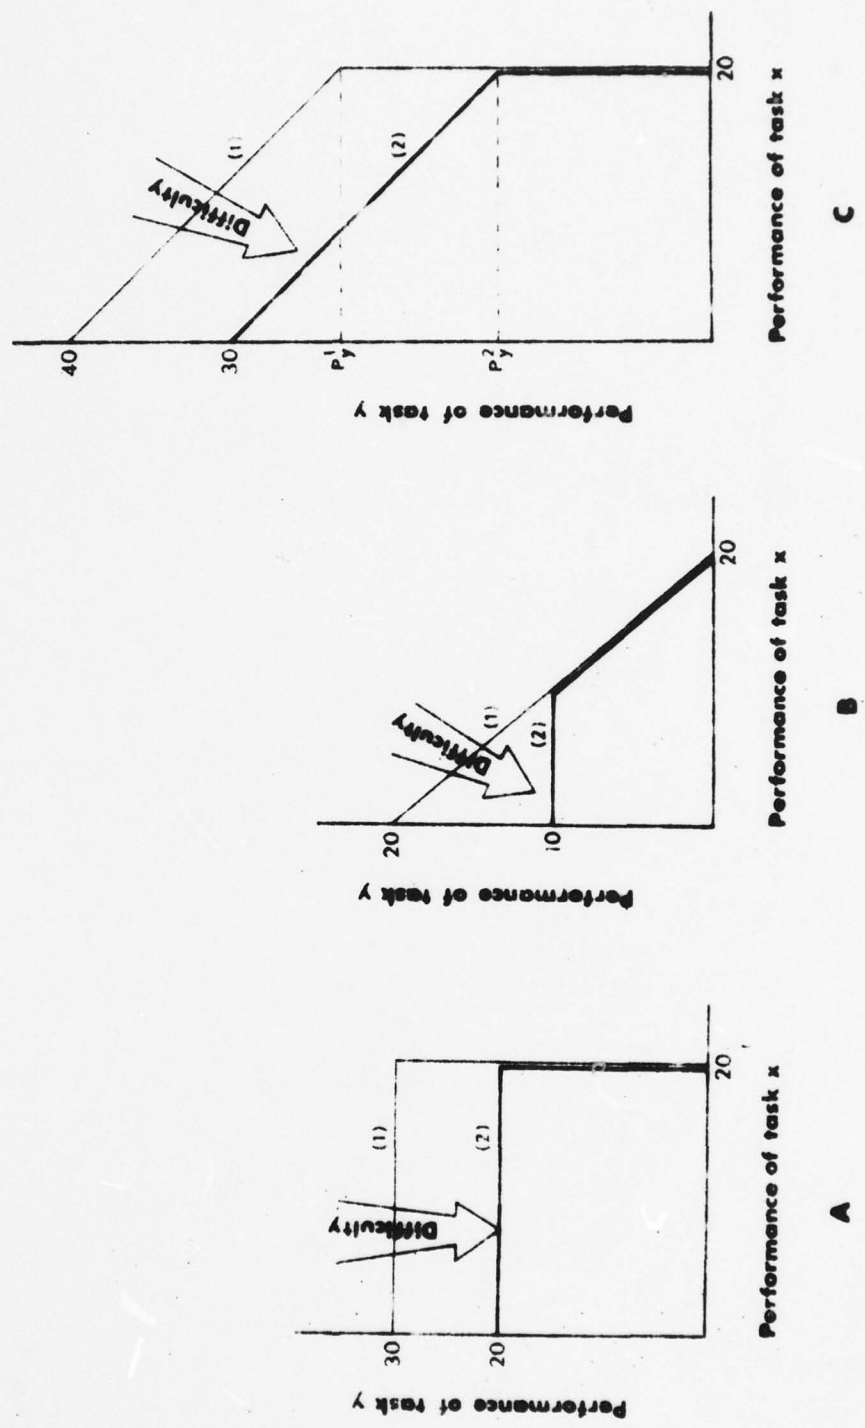


Fig.4: Illustrations for three possible effects of manipulating a parameter which taps on resources relevant only for the performance of y. Panel A presents the case of complete absence of common resources. Panel B presents the case of existence of both common and disjoint resources used in fixed proportions. Panel C presents the case of common and disjoint resources used substitutively.

of R^2 . The corresponding POC is curve 1 in Figure 4B. Now suppose that a parameter of task y is manipulated so as to increase its demand for R^2 to two units. This sets a new limit to the performance of y which is 10, but since to achieve that level the performer needs no more than 10 units of R^1 , the rest can be directed to task x. The resulting POC is curve 2 in Figure 4B.

Now let us examine what happens when we relax the requirement that resources will be employed in fixed proportions. Suppose that there is more than one way to do a task. There may be one optimal composition of resources, but deviations are tolerated and performance usually benefits to some extent from increases of one type of resources, even when not accompanied by commensurate increases of other types. For instance, although usually two units of STM are used along with one unit of VIS, the process makes some use of a third unit of STM even when one unit of VIS is available. In this case, the types of performance functions are illustrated in Figure 5B and 5C by means of iso-performance contours as a function of two types of resource. The extreme case of variable pro-

Insert Figure 5 about here

portions illustrated in Figure 5C is when the use of the two types of resource is perfectly substitutive. More typical presumably is the case of partial substitution illustrated in Figure 5B in which substituting R^1 for R^2 (or vice versa) is progressively less productive. As a reference point consider the illustration of the fixed proportions case in Figure 5A: the ratio of three units of R^1 to one unit of R^2 is mandatory.

20A

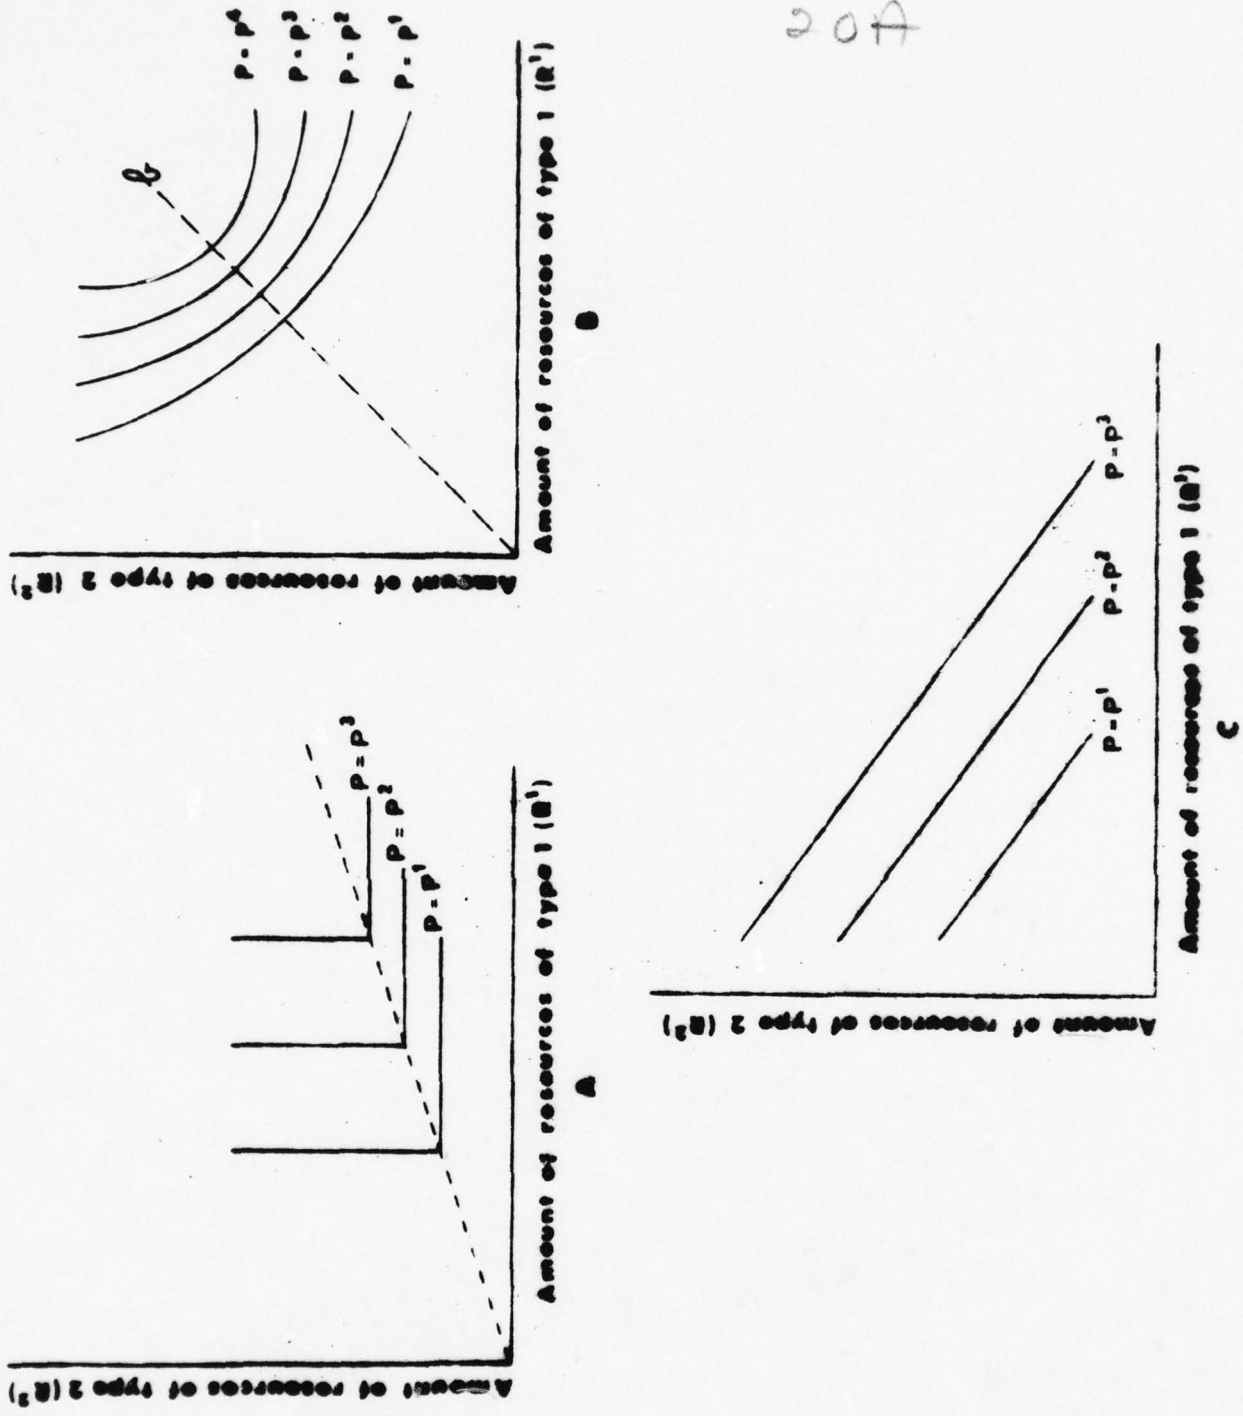


Fig.5: Iso-performance contours as a function of two types of resource, R¹ and R². Each contour connects all resource-combinations that yield the same level of performance (e.g., P²). Panel A presents a fixed-proportions performance function. Panels B and C present variable-proportions performance functions.

When the proportions of input resources are not fixed, the performer can reduce interference between concurrent tasks by operating less with the common resources and more with the disjoint ones. In the extreme case of variable-proportions the use of different types of resource is perfectly disjunctive, namely either of several resources can be employed by itself (see Figure 5C). Take the example illustrated in Figure 4B and just change the italicized and to or, and you will obtain the effect demonstrated in Figure 4C: a reduction in the efficiency of the disjoint resource R^2 decreases the level of performance of y which can be attained with no cost to the performance of x (see the decrease from P_y^1 to P_y^2 in Figure 4C). If the use of resources is just partly substitutive (see Figure 5B), then to maintain the performance of y at the same level, a larger amount of R^1 has to be invested in it, thus task x suffers. It is hard to predict whether it will suffer more or less with higher performance of y . If the iso-performance contours are progressively more concave (as in Figure 5B), it will suffer less; if they are equally concave, it will suffer the same, so the family of POCs will look as in Figure 4C; if they are progressively less concave, task x will suffer more the higher the performance of task y . In this case a manipulation which taps a disjoint resource will produce a fan-like family of bowed-out POCs that is practically indistinguishable from the effect of manipulating the demand for a common resource.

This analysis should call our attention to an important point: the overlap in cost composition of concurrent tasks may be partial, that is they may use some common resources and at the same time each task may resort to some resources not required by the other one. Hence, a failure of a manipulation of the difficulty of one task to affect the performance

of the other one (when the performance of the first one is held constant) just proves that resource overlap is not total but not that it does not exist. On the other hand, even when joint performance exhibits a considerable trade-off due to shifts in resource allocation, the tasks may still depend on some different mechanisms which can be detected by manipulating various subject-task parameters and observing effects like the ones in Figures 4B and 4C. The lesson is that there is one more good reason for researchers to do what we already recommended, namely to manipulate subject-task parameters as well as task preferences and to present their effects in terms of families of POCs. Different parameters may yield different pictures depending on the resources which they tap on.

We employed this approach in a study of the interaction between axes in two-dimensional pursuit tracking (Gopher & Navon, in preparation). We regarded this situation as time-sharing between horizontal and vertical tracking and measured tracking error in each of the dimensions. We controlled relative emphasis on the two dimensions by varying the ratio of tolerance levels for error in each, and manipulated the difficulty of each dimension independently by varying some parameter of tracking in that dimension. In the first experiment the manipulated variable was the cut-off frequency for the low-pass filter applied to the output of a random noise generator to yield the target forcing function. In the second experiment difficulty was manipulated by changing the target velocity (which was also higher on the average than the velocity in the other experiments). In the third experiment we varied the ratio of acceleration to velocity in the control dynamics of the hand controller (this ratio was also higher on the average than in the other two experiments). Some of the results are presented in Figure 6 as families of POCs.

In the first experiment (see Figure 6A) task emphasis had a large

Insert Figure 6 about here

effect on performance (a range of about 13 percent of scale Root Mean Square tracking error), which was nevertheless negatively accelerated as indicated by the strong curvature of the POCs: subjects did respond to a lower requirement in one dimension by increasing tracking error, but that helped them very little to improve performance in the other dimension. The difficulty effects were much smaller: the frequency of the target vertical movement affected vertical accuracy linearly but horizontal accuracy curvilinearly, and it did not interact with the task emphasis variable⁷.

The results of the second and the third experiments (see Figures 6B and 6C respectively) are characterized by a smaller effect of task emphasis and a larger effect of the difficulty manipulations. But here the similarity ends. The limited performance trade-off exhibited in the first experiment recurs in the second one. The velocity manipulation had a linear effect just on the manipulated axis and did not interact with task emphasis.

In contrast, the fan-like family of almost linear POCs in Figure 6C reveals that in the third experiment performance tradeoff was clear and that the manipulation of the control dynamics interacted with task emphasis⁸.

The results can be accommodated nicely within a post-hoc account based on the notion of multiple resources. Suppose that despite the apparent similarity between vertical and horizontal tracking, their cost compositions are fairly disjoint. Suppose further that the two tasks require the same kind of motor-related resource but different kinds of perceptual or "computational" resources. Now, in the first two experiments the load on the motor system, which was the common resource, was



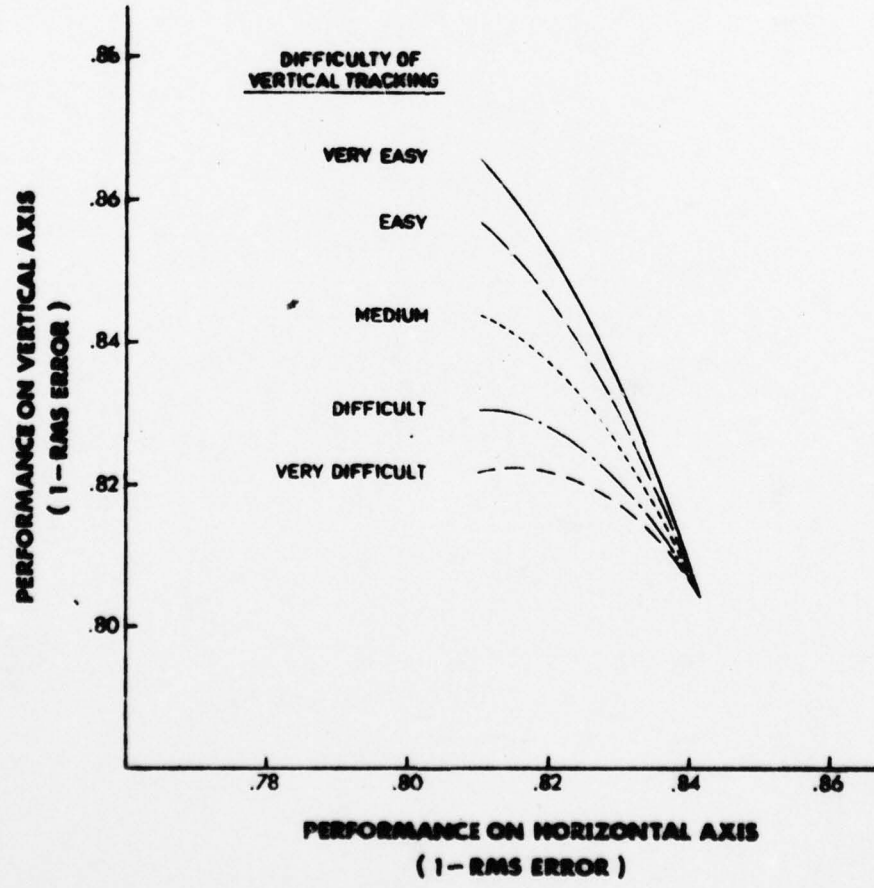


Fig.6: A family of POCs representing tracking accuracy (1-Root Mean Square Error) on each of the axes in dual-axis tracking as a function of task emphases. Each POC corresponds to a different level of difficulty of vertical tracking, and is obtained by jointly solving two second-order multiple regression equations for predicting performance on the two axes from the task emphases variable. The different panels correspond to different difficulty manipulations: frequency of target movement in panel A, target velocity in panel B, and control dynamics in panel C. Panel D presents average performance combinations in each level of task emphasis in the second experiment with the POC fitted to them.




B



C

-23C-



relatively small, hence the tasks did not interfere very much with each other. However, because of the higher average ratio of acceleration to velocity in the third experiment, both tasks required more motor capacity, so that they had more to compete for. The different effects of the difficulty manipulations may be interpreted if we realize that the only parameter which seems to affect the motor system is the control dynamics manipulation in the third experiment. Manipulations which tap a common resource are expected to yield a fan-like family of POCs (see Figure 3), whereas ones which tap disjoint resources may yield a family like the one illustrated in Figure 5C (which is basically what was found in the second experiment). The curvilinearity of the effect of the frequency of vertical movement on horizontal accuracy (Figure 6A) is hard to interpret. The lower variance due to the task emphasis variable in the last two experiments, which were on the average more difficult than the first one, is probably a manifestation of the smaller effect of investment of resources when their efficiency is low, that is when the task is difficult (see Figure 1).

However, if these two tasks are not similar enough to call for exactly the same resources, one might wonder whether the number of different resources identified empirically will not turn out to be too large to make the notion of multiple resources useful. Thus, this analysis illustrates both the potential utility of interpretations in terms of multiple resources as well as the difficulties.

Resources and Stages

What do we mean when we say that different resources collaborate in processing? What do we mean when we say that tasks are time-shared? Different processes may operate in parallel, in sequence, or intermittently, and it is quite difficult to diagnose experimentally which mode of time-sharing is actually taking place in a given situation (see Townsend, 1974).

We do not address these issues here, but we would like to present just two implications from our analysis to the application of additive factors logic (Sternberg, 1969) in associating tasks or factors with processing stages.

First, Sternberg proposed that the effect on latency of two factors affecting the same stage should typically be interactive. This idea which has been employed in numerous studies, was originally put forward without much theoretical justification (e.g., Sternberg, 1969, p.282), except for the intuitively appealing cognitive symmetry it creates with Sternberg's other suggestion that the effects of factors affecting different stages should be additive. However, the validity of this claim seems dubious, if we construe the effects of experimental factors as changes in resource efficiency. To see this it is best to consider economic analogies. For example, the marginal output of labor may vary as a function of, say, the skill of laborers and their industriousness. These attributes may or may not interact. The same thing holds for factors influencing different resources: whether or not the skill of laborers and the quality of instruments interact is an empirical question. To predict the composition rule of factor effects one needs to know about the process of production much more than just that "both factors affect the same stage". Hence, interaction may be sufficient to rule out the hypothesis that factors influence different stages, yet additivity is not sufficient for rejecting the alternative hypothesis (cf. Taylor, 1976).

Second, the additive factors logic is sometimes married with capacity models to account for phenomena of dual-task performance. The

rationale here is that an interaction between a certain factor and a dual-task requirement (namely, the presence of differential effects of that factor on latency to perform the task in single- and dual-task situations) indicates that that factor and the concurrent task affect the same stage, whereas additivity suggests an influence upon different stages (see, e.g., Griffith & Johnston, 1977). However, this application of the additive factors logic seems even more problematic. Interaction of this sort is possible even when the manipulated factor affects a stage during which the concurrent tasks do not compete for capacity, provided that they do compete at some other stage. For example, suppose a subject is required to name a digit and concurrently to manually respond to its mere appearance, and suppose we test the effect of backward masking on the manual detection latency (see experiment reported by Kantowitz, in press). Let us assume that masking influences just the encoding stage of the naming task, and that naming and detection do not compete for capacity during encoding (perhaps because encoding does not require any capacity, see Posner & Boies, 1971). At first glance, it appears that detection latency should be inhibited equally by naming either a masked or a nonmasked digit. However, this is true only if subjects let masking affect just naming latency. If they nevertheless aim at protecting the naming task, they may try to compensate for the inhibitory effect of masking on the encoding stage by increasing the share of capacity directed to the naming task during response selection. That may result in a delay of the detection response which is, thus, a capacity interference indirectly caused by the apparently irrelevant mask⁹. This is another example for the importance of having control over, or at least knowledge about, task preferences before drawing any conclusions about processing.

Summary

The effect of task difficulty on performance can be conceptualized within a theory which posits that performance depends on the use of resources from a single pool. When the difficulty of a task is said to increase, it may mean either that resources invested in it can now do less (i.e. a decrease in efficiency) or are now required to do more (i.e. an increase in load), or have now less time to do it (i.e. a stricter limit on processing duration). Either way, difficulty should most often interact with resource investment in such a way that effects of resource investment on quality or speed of performance are more pronounced the easier the task is.

If the processing system is viewed as comprised of a number of mechanisms each having its own capacity, then a difficulty manipulation may affect differentially the use of each of those capacities. If in a dual-task situation a manipulation of the difficulty of one task affects the use of a mechanism which is not required by the other task, processing of the latter may remain intact under some circumstances.

To get a complete picture of how difficulty affects dual-task performance, it is proposed to manipulate task preferences as well as difficulty parameters and present their joint effect by families of POCs.

REFERENCES

- Briggs, G.E., Peters, G.L., & Fisher, R.P. On the locus of divided attention effects. Perception and Psychophysics, 1972, 11, 315-320.
- Broadbent, D.E. Decision and Stress. London: Academic Press, 1971.
- Gopher, D., & North, R.A. The measurement of attention capacity through concurrent task performance with individual difficulty levels and shifting priorities, Proceedings of the Eighteenth Annual Meeting of the Human Factors Society, Santa Monica, California, October 1974.
- Griffith, D., & Johnston, W.A. Stage 2 processing and the divided-attention effect. Memory and Cognition, 1977, 5, 597-601.
- Kantowitz, B.H. Channels and stages in human information processing: A limited review. Cognitive Psychology, in press.
- Kantowitz, B.H., & Knight, J.L., Jr. Testing tapping time-sharing, Journal of Experimental Psychology, 1974, 103, 331-336.
- Kantowitz, B.H., & Knight, J.L., Jr. Testing tapping time-sharing, II: Auditory secondary task. Acta Psychologica, 1976, 40, 343-362.
- Kantowitz, B.H., & Knight, J.L., Jr. When is an easy task difficult and vice versa? Acta Psychologica, in press.
- Kerr, B. Processing demands during mental operations, Memory and Cognition, 1973, 1, 401-412.
- Lane, D.M. Attention allocation and the relationship between primary and secondary task difficulty. A reply to Kantowitz and Knight. Acta Psychologica, 1977, 41, 493-495.

- Moray, N. Where is capacity limited? A survey and a model, Acta Psychologica, 1967, 27, 84-92.
- Navon, D., & Gopher, D. On the economy of the human processing system. In preparation.
- Norman, D.A., & Bobrow, D.J. On data-limited and resource-limited processes, Cognitive Psychology, 1975, 7, 44-64.
- Norman, D.A., & Bobrow, D.J. On the analysis of performance operating characteristics, Psychological Review, 1976, 83, 508-510.
- Posner, M.I., & Boies, S.W. Components of attention, Psychological Review, 1971, 78, 391-408.
- Rolfe, J.M. The secondary task as a measure of mental load. In W.T. Singleton, R.S. Easterby and D.E. Whitfield (Eds.), Measurement of Man at Work, London: Taylor & Francis, 1971.
- Shiffrin, R.M. Capacity limitations in information processing, attention, and memory. In W.K. Estes (Ed.), Handbook of Learning and Cognitive Processes (Vol.4), Hillsdale, N.J.: Lawrence Erlbaum, 1976.
- Sternberg, S. The discovery of processing stages: Extension of Donders' method. In W.G. Koster (Ed.), Attention and Performance II, Amsterdam: North-Holland, 1969.
- Taylor, D.A. Stage analysis of reaction time, Psychological Bulletin, 1976, 83, 161-191.
- Townsend, J.T. Issues and models concerning the processing of a finite number of inputs. In B.H. Kantowitz (Ed.), Human Information Processing: Tutorials in Performance and Cognition, Hillsdale, N.J.: Lawrence Erlbaum, 1974.

Wattenbarger, B.L., & Pachella, R.G. The effect of memory load on reaction time in character classification. Perception and Psychophysics, 1972, 12, 100-102.

FOOTNOTES

2. If the criterial output reflects also a ceiling of performance, as the case is, for instance, with accuracy measures, then adequacy = $\text{MIN}(1.00, \text{processing duration} \times \text{output rate/load})$. Note, however, that the criterion may be sometimes defined in such a way that actual output can exceed criterial output, for example when there exists some nonzero tolerance range in tracking.
3. This assumes, of course, no guessing.
4. Note that one can treat load/processing-time as criterial rate. While this merge of variables does not seem very meaningful in discrete tasks such as visual search, it may be more natural in continuous tasks such as tracking.
5. Of course, processing duration does not necessarily equal SOA, but it is presumably highly correlated with it.
6. It is possible to distinguish between marginal and average cost, but we will not pursue this distinction here. To simplify analysis, we assume that all marginal cost compositions preserve the same proportions of specific resources.
7. We do not present in Figure 6 the effects of the difficulty manipulations on the horizontal axis, but they are basically similar.

8. To give an idea of the goodness of fit of the curves to the data, the actual performance combinations in the second experiment are presented in Figure 6D averaged across subjects and levels of difficulty for each of the levels of task emphasis, along with the POC curve fitted to them. The fit is slightly better in the first experiment and slightly worse in the third one.
9. An alternative mechanism that may produce the same effect and that is not related to capacity interference but rather to response strategies was conjectured by Kantowitz (in press).

Unclassified

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. Report Number AFOSR-TR-79-0828	2. Govt Accession No.	3. Recipient's Catalog Number
4. Title (and Subtitle) Interpretations of task difficulty in terms of resources: Efficiency, load, demand, and cost composition.	5. Type of Report & Period Covered Final report Oct. 1, 1977 - Sept. 30, 1978	
	6. Performing Org. Report Number 78-1 ✓	
7. Author(s) David Navon, Daniel Gopher	8. Contract or Grant Number AFOSR-78-3131 <i>new</i>	
9. Performing Organization Name and Address Technion-Israel Institute of Technology, Technion City, Faculty of Industrial & Management Engineering, Haifa, Israel.	10. Program Element, Project, Task Area & Work Unit Numbers 61102F 2313 A2	
11. Controlling Office Name and Address Life Sciences Directorate (NL) Air Force Office of Scientific Research Bolling Air Force Base, Washington, D.C.20332.	12. Report Date November 1978	
	13. Number of Pages 31	
14. Monitoring Agency Name and Address	15.	
16. & 17. Distribution Statement Approved for public release; distribution unlimited.		
18. Supplementary Notes		
19. Key Words Difficulty, capacity, load, workload, time-sharing performance, task priority, task interference, task demand, performance tradeoff, processing resources		
20. Abstract The effect of task difficulty on performance can be conceptualized within a theory which posits that performance depends on the use of resources from a single pool. When the difficulty of a task is said to increase, it may mean either that resources invested in it can now do less (i.e. a decrease in <u>efficiency</u>) or are now required to do more (i.e. an increase in <u>load</u>), or have now less time to do it (i.e. a stricter limit on <u>processing duration</u>). Either way, difficulty should most often interact with resource investment in such a way that effects of resource investment on quality or speed of performance are more pronounced the easier the task is. (cont..)		

Unclassified

9 *Unclassified*

Abstract (Cont.)

If the processing system is viewed as comprised of a number of mechanisms each having its own capacity, which may be considered as a separate *resource*, then a difficulty manipulation may affect differentially the use of each of those capacities. If in a dual-task situation a manipulation of the difficulty of one task affects the use of a mechanism which is not required by the other task, processing of the latter may remain intact under some circumstances.

To get a complete picture of how difficulty affects dual-task performance, it is proposed to manipulate task preferences as well as difficulty parameters and to present their joint effect by families of POCs. An application of this methodology to the study of pursuit tracking is briefly described and interpreted in terms of multiple resources.

Unclassified