Voluntary Control of Human Cardiovascular Integration and Differentiation through Feedback and Reward

Abstract. *Human subjects can learn to control the relation between their systolic blood pressure and heart rate when they are given feedback and reward for the desired pattern of blood pressure and heart rate. They can learn to integrate these functions (increase or decrease both jointly), or to a lesser degree, differentiate them (raise one and simultaneously lower the other). The extent of this learning is predicted by a behavioral and biological model that explains specificity of learning in the autonomic nervous system.*

The traditional notion of the autonomic nervous system is of a tightly controlled homeostatic network capable of little response separation or voluntary control. Contrary to this belief, recent research on instrumental learning of visceral responses has shown that individual functions can be brought under voluntary control and can show specificity of learning similar to that found for skeletal responses (*1*). For example, research on systolic blood pressure (BP) and heart rate (HR), two closely related autonomic responses, has shown that each can be separately controlled. (i) Providing human subjects with feedback and reward for increases or decreases in systolic BP leads to learned control of BP without corresponding changes in HR (*2, 3*). (ii) If increases or decreases in HR are reinforced, the subjects learn to control HR without similarly changing BP (*4*). Although this phenomenon has both theoretical and clinical importance (for example, as a potential treatment for decreasing specific symptoms in psychosomatic disorders), a suitable explanation for it has not yet appeared.

The purpose of the proposed integration-differentiation (ID) model (*5*) is to help provide a behavioral and biological framework for understanding and predicting learned patterns of physiological activity. The model takes into account (i) the behavioral relations between responses as defined by the operant (feedback) procedure (*6*) and (ii) natural physiological changes or constraints that occur over time. In this context, the term integration is reserved for the response pattern in which two functions simultaneously change in the same direction (both increasing and decreasing together in a sympathetic-like pattern), and the term differentiation refers to the response pattern in which two functions simultaneously change in opposite directions (*7*). The first part of the model is an attempt to assess the activity of other physiological functions at the instant when a given physiological function is being reinforced, and the second part is an attempt to relate this information to naturally occurring changes in physiological activity due to stimulation, adaptation, and other factors not related to the contingency of reinforcement per se (*8*).

The general approach can be illustrated through a behavioral analysis of systolic BP and HR, two functions showing discrete bursts of activity that can be easily reduced to binary response units. With every heart contraction, BP rises and reaches a systolic peak, its magnitude determined by a complex interaction of HR, stroke volume, and peripheral resistance (*9*). If average (tonic) levels are specified for HR (in beats per minute) and for BP (in millimeters of mercury), each heart beat can be classified according to whether HR and BP values are above $(^{up})$ or below $(_{down})$ their tonic levels. At each heart beat only four coincidence patterns are possible: BP^{up} HR^{up}, BP_{down} HR_{down}, BP^{up} HR_{down}, and BP_{down} HR^{up}.

If the two functions naturally rose and fell together all of the time (BP^{up} HR^{up} and BP_{down} HR_{down} only), then even though an experimenter might select only one system for reward, he would unwittingly provide the same reinforcement for parallel changes in the other system. Accordingly, both functions would show simultaneous learning in the same direction (integration). At the other extreme, if the two functions always changed in opposite directions (BP^{up} HR_{down} and BP_{down} HR^{up} only), then when changes in one function were selected for reward, opposite changes in the other function would also receive contingent (100 percent correct) reinforcement. Again, both functions would show simultaneous learning, this time in opposite directions (differentiation). However, if the two functions were unrelated to the point of actually producing equal numbers of the four coincidence patterns, then when changes in one function were reinforced, the other function would receive a sequence of random reinforcement. (For example, if every BP^{up} was rewarded, HR would simultaneously be rewarded half the time for HR^{up} and half for HR_{down}.) Therefore, only the response receiving contingent reinforcement would be learned. I refer to this latter situation (when only the response chosen for reinforcement is learned) as specificity of learning.

My experiment was designed with two major purposes. One was to empirically determine the naturally occurring behavioral relationship between BP and HR, to test the hypothesis that, in order for specificity of BP and HR learning to occur, these two functions vary independently and thus go up and down together only about 50 percent of the time. The second aim was to answer the following question: if HR and BP (or any two functions) are randomly related, then is it possible to make them change together, or change in opposite directions? One implication of the ID model is that it should be possible to control a combination of functions, for example BP and HR, by rewarding the subject only when he shows the desired coincidence pattern of simultaneous changes in both functions (*10*). This procedure, although it reduces the percentage of correct reinforcement each function can receive (since neither function alone controls the reward), eliminates incorrect (undesired) rewards. A behavioral analysis would predict that if the frequency of the four coincidence patterns of BP and HR were equal and randomly distributed, then each pattern would show learning to the same degree. However, this prediction would fail if physiological constraints were preventing HR and BP from changing independently. In this case, the discrepancy between the predicted and measured extents of learning would uncover the ways in which the constraints were operating.

This approach requires a method for measuring on-line the phasic (beat-by-beat) and tonic (median) relationships between autonomic responses. Instrumentation was developed which determines at each heart beat the four coincidence patterns for BP and HR and also provides tonic value for each function during each trial (*11*) (Fig. 1).

Forty normotensive males (paid volunteers), 21 to 30 years old, were seated in a sound- and temperature-controlled room and connected to the physiological recording devices. Systolic BP, HR, and respiration were recorded on an Offner type R dynagraph. The electrocardiogram was measured with standard plate electrodes and displayed on one channel of the dynagraph. An electronic switch triggered Grason-Stadler (model 1200) solid-state programming equipment at each heart cycle (R spike of the electrocardiogram). The equipment automatically detected the four coincidence patterns for BP and HR and also presented all stimuli to subjects. Respiration was recorded by a strain gauge placed around the waist.

For each trial, median systolic BP was defined as the constant cuff pressure at which 50 percent of possible Korotkoff sounds occurred. Korotkoff sounds were displayed on one channel of the dynagraph in conjunction with a second electronic switch. With the use of this switch and of appropriate logic modules it was possible to count heart beats accompanied (within 300 msec) by a Korotkoff sound and heart beats not followed by a Korotkoff sound. Similarly, median HR was obtained by the use of a third electronic switch calibrated in beats per minute. The HR for each beat was displayed on one channel of the dynagraph through a cardiotachometer (Lexington Instruments model 107). With the use of appropriate logic modules in conjunction with the third (cardiotachometer) electronic switch, it was possible to count heart beats that were faster or slower than the value set by the cardiotachometer electronic switch.

If Korotkoff sounds (BP$^{\text{up}}$) or fast heart beats (HR$^{\text{up}}$) or both exceeded the median levels by 35 in a given trial (50 heart beats in length), on the next trial the applied cuff pressure was raised 2 mm-Hg, or the cardiotachometer electronic switch was raised 2 beats per minute, or both changes were made. If the number of BP$^{\text{up}}$ or HR$^{\text{up}}$ or both exceeded the medians by less than 15 in a given trial, on the next trial the applied cuff pressure was lowered 2 mm-Hg, or the cardiotachometer electronic switch was lowered 2 beats per minute, or both changes were made. This shaping procedure made it possible to accurately track both BP and HR independently and simultaneously, and at the same time to obtain comparable information about relative changes in BP and HR at each beat.

Subjects were told that the purpose of the experiment was to determine whether they could learn to control certain physiological responses that are generally considered involuntary. They were instructed to refrain from moving and to breathe regularly. They were not told the nature of the specific responses to be controlled or the required directions of change. Twenty subjects were studied in two integration conditions. Half of these subjects received a 100-msec light and tone as feedback for each pulse cycle in which BP and HR simultaneously increased (BP$^{\text{up}}$ HR$^{\text{up}}$); the other half received the same feedback each time their BP and HR simultaneously decreased (BP$_{\text{down}}$ HR$_{\text{down}}$). Another 20 subjects were studied in two differentiation conditions. Half received the feedback only when their BP increased and their HR simultaneously decreased (BP$^{\text{up}}$ HR$_{\text{down}}$), and the other half received the feedback for BP$_{\text{down}}$ HR$^{\text{up}}$ responses.

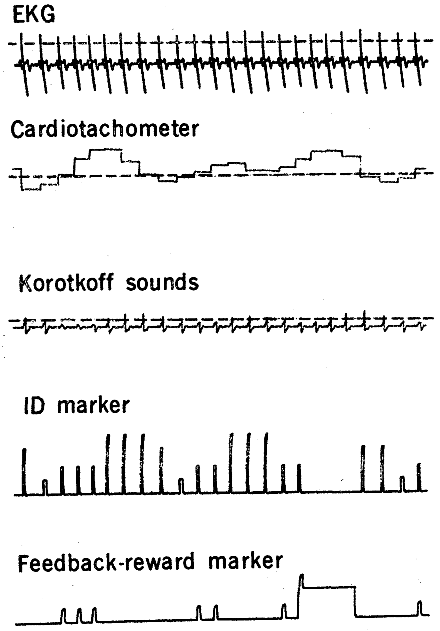

Fig. 1. Representative portion of a polygraph record of the integration-differentiation (ID) system in operation. Shown are the electrocardiogram (EKG), heart rate (HR) displayed through a cardiotachometer, Korotkoff sounds measured at a constant cuff pressure, and two marker channels. Dashed lines represent the approximate levels of the three electronic switches. The presence or absence of a Korotkoff sound relative to the constant pressure in the cuff indicates whether blood pressure (BP) is up or down, while HR is rated up or down relative to the median HR. After each heart cycle (except during a reward) one of four possible marks appears on the ID marker channel. The longest and shortest marks indicate integration, with BP$^{\text{up}}$ HR$^{\text{up}}$ producing the longest mark and BP$_{\text{down}}$ HR$_{\text{down}}$ producing the shortest mark. The other two marks indicate differentiation, with BP$^{\text{up}}$ HR$_{\text{down}}$ producing the third longest mark, and BP$_{\text{down}}$ HR$^{\text{up}}$ producing the second shortest mark. The bottom channel indicates which one of the four possible combinations is eliciting feedback (short mark) and reward (long mark). In this example, feedback is occurring for BP$_{\text{down}}$ HR$^{\text{up}}$ differentiation.

After giving every 12 correct responses, subjects were rewarded with a 3-second view of slides showing landscapes, attractive nude females, or cumulative bonuses earned (each slide was worth 5 cents). No feedback was given during the slides. All subjects receive 5 adaptation, 5 random reinforcement, and 35 conditioning trials. A blue light signaled the onset of random reinforcement and conditioning trials and remained on during the trials. Each trial was 50 heart beats long and was preceded by 10 seconds of cuff inflation. Intervals between trials ranged from 20 to 30 seconds; during this time the cuff was deflated. Subjects were matched according to their resting median BP's and randomly placed in one of the four conditions (*12*).

Figure 2 shows the average conditioning results for the four groups, based on the median cuff pressures and HR's obtained in each trial. Results are summarized as follows.

1) Subjects learned to directly integrate their BP and HR in a single session. Separate analyses of variance revealed highly significant group-by-trial interactions for BP ($P < .0001$) and HR ($P < .0001$), a result which indicates that the divergence of the two BP curves and of the two HR curves is reliable (*13*). Note that when BP$^{\text{up}}$ HR$^{\text{up}}$ is rewarded, both BP and HR increase somewhat and then return to baseline, while BP$_{\text{down}}$ HR$_{\text{down}}$ reinforcement yields sustained decreases in both. This result is very similar to earlier data obtained when each was separately rewarded (*2*, *4*), except that in the present experiment control of both BP and HR is learned, and the learning effect is greater.

2) Integration was learned more easily than was differentiation. In fact, each of the integration groups earned more slides than either differentiation group. The divergence of the two BP curves in Fig. 2 (right) is reliable ($P < .0001$) and is very similar to the results for integration. In contrast, the divergence of the two HR curves in Fig. 2 (right) is not reliable, primarily because of the lack of sustained HR de-

creases in the $\mathrm{BP^{up}\ HR_{down}}$ group (*13*).

3) If the curves in Fig. 2 (left and right) are compared, the $\mathrm{BP^{up}\ HR_{down}}$ group is similar to the $\mathrm{BP^{up}\ HR^{up}}$ group, while the $\mathrm{BP_{down}\ HR^{up}}$ group is more similar to the $\mathrm{BP_{down}\ HR_{down}}$ group.

4) Clear evidence for learning of differentiation was obtained only for the $\mathrm{BP_{down}\ HR^{up}}$ group ($P < .0001$); the reverse pattern ($\mathrm{BP^{up}\ HR_{down}}$) proved difficult to learn (*14*).

These results were anticipated from the ID model, by an analysis of the natural beat-by-beat and tonic relation of BP to HR during rest and during random stimulation in the present experiment and by a comparison of these data with earlier data on BP and HR control (*2–4*). A summary of the predictions from the model follows.

1) The $\mathrm{BP^{up}\ HR^{up}}$ pattern would look as if it were learned less effectively than the reverse $\mathrm{BP_{down}\ HR_{down}}$ pattern, when the levels found at random stimulation were used as the standard. This prediction was based on data suggesting that systolic BP and HR tend to normally adapt (decrease relative to initial rest or random stimulation levels) during this type of experiment (*2*). Adaptation would therefore act to lower both the increase and decrease curves for BP and HR.

2) Tonic integration would be more readily learned than tonic differentiation. This prediction was based on the observation that median levels of BP and HR tend to increase (or decrease) together ($r = +.36$, $P < .05$) during the initial period of random stimulation. (For example, if a subject reacts to the beginning reinforcement trials with a large tonic increase in BP, his HR will tend to show a comparable rise in level.) In other words, to the extent that the systems are tonically integrated in this experimental situation (other situations might heighten, eliminate, or possibly reverse this relationship), it should be easier to make them change together than in opposite directions.

3) Subjects reinforced for the $\mathrm{BP^{up}\ HR_{down}}$ pattern would raise both functions (and thus appear similar to the $\mathrm{BP^{up}\ HR^{up}}$ group); while subjects reinforced for the $\mathrm{BP_{down}\ HR^{up}}$ pattern would tend to lower both functions. This hypothesis is based on the following analysis of the BP and HR coincidence patterns. As predicted from the specificity findings (*2–4*), BP and HR change in the same direction only 53 percent of the time, and the frequency of the four coincidence patterns is approximately equal. However, close inspection of the coincidence patterns reveals that the relation between BP and HR over time is not truly random. Changes in HR precede those of BP; this relation produces a constant sequence of coincidence patterns (*15*). For example, the data indicate that when a $\mathrm{BP^{up}\ HR^{up}}$ pattern changes, 70 percent of the changes will be to a $\mathrm{BP^{up}\ HR_{down}}$ pattern. Therefore, to the extent that the $\mathrm{BP^{up}\ HR^{up}}$ pattern consistently precedes the $\mathrm{BP^{up}\ HR_{down}}$ pattern, then reinforcement for the latter will produce consistent (although delayed) reinforcement of the former as well. This is an example of a response chaining factor. Similarly, the data indicate that when a $\mathrm{BP_{down}\ HR_{down}}$ pattern changes, it will change with a 70 percent frequency to a $\mathrm{BP_{down}\ HR^{up}}$ pattern; this observation explains the similarity to $\mathrm{BP_{down}\ HR_{down}}$ learning when $\mathrm{BP_{down}\ HR^{up}}$ responses are reinforced. However, with natural tonic integration acting as a constraint, the response chaining factor may have little influence when integration patterns for BP and HR (both increasing and decreasing) are directly reinforced.

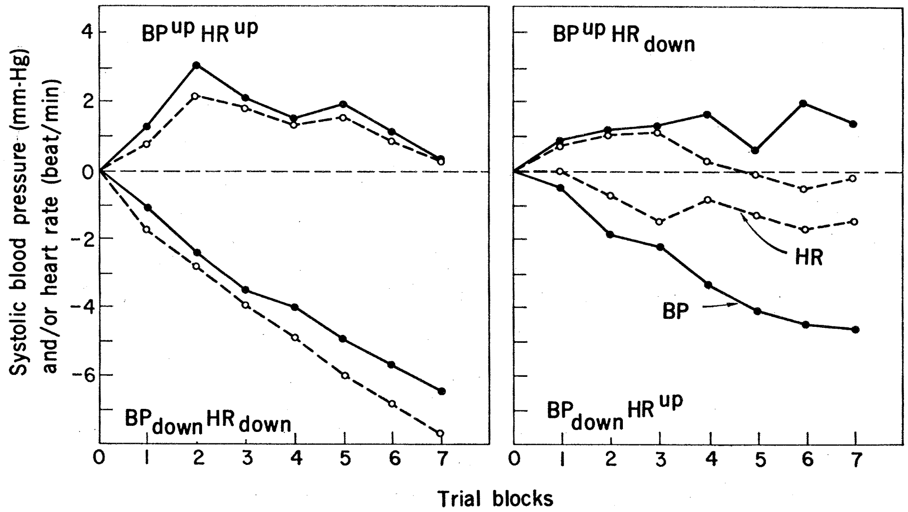

Fig. 2. Average systolic blood pressure (BP) and heart rate (HR) for the subjects being conditioned in the four coincidence patterns. On the left are data for the two integration conditions, on the right are data for the two differentiation conditions. Solid lines are BP, dashed lines are HR. Each point is the mean of five trials, set to zero by the last random trial. Beats per minute and millimeters of mercury are therefore on the same axis.

4) The $\mathrm{BP_{down}\ HR^{up}}$ pattern would be more readily learned than the opposite $\mathrm{BP^{up}\ HR_{down}}$ pattern. This was predicted from the interaction of tonic and phasic factors in the present reinforcement procedure. To the extent that BP and HR normally adapt within a trial (but the electronic switches remain constant), the frequency of $\mathrm{BP_{down}\ HR_{down}}$ patterns would tend to increase toward the end of a trial. The $\mathrm{BP_{down}\ HR^{up}}$ pattern naturally occurs after a $\mathrm{BP_{down}\ HR_{down}}$ pattern, whereas the $\mathrm{BP^{up}\ HR_{down}}$ pattern rarely occurs after a $\mathrm{BP_{down}\ HR_{down}}$ pattern. In other words, it appears more feasible to change from a $\mathrm{BP_{down}\ HR_{down}}$ pattern by raising HR than by raising BP (*16*).

To test for conditioning of respiration during the experiment, analyses were made of respiration rate, variation of respiration rate, and variation of respiration amplitude. No differences were found between the four groups of subjects (*17*). Also, cognitive activity was assessed at the end of the experiment by a comprehensive questionnaire. Subjects in the $\mathrm{BP_{down}\ HR_{down}}$ group tended to check more items associated with relaxation than did the $\mathrm{BP^{up}\ HR^{up}}$ group, a result suggesting that some kind of arousal variable may have been operating. However, the two differentiation groups were indistinguishable from the $\mathrm{BP^{up}\ HR^{up}}$ group. Apparently, when a person is required to decrease the activity of more than one function simultaneously, subjective relaxation may occur.

Altogether, these data demonstrate that it is possible for human subjects to develop some control over the relation between their systolic BP and HR when they are provided with feedback and reward for the desired pattern of BP and HR. When the behavioral relation of these functions as "seen through the eyes of feedback and reward" is analyzed, and when natural biological constraints are taken into account, it appears possible to understand and predict the resulting pattern of learned control (*18*). To this extent, these results clarify earlier specificity findings and provide a new framework for research and theory in the control of multiautonomic functions. It is suggested that the present technique may be a tool for studying and controlling not only the relationship between visceral responses but the interaction of visceral responses with somatic and central behavior as well (*19*). This could be accomplished by assessing the extent and ease with which specific patterns of activity can be learned. The importance of considering feedback and reward in biological perspective is stressed, since natural physiological relationships and constraints do occur (*20*). It may be possible to apply these techniques to the treatment of specific clinical disorders; for example, to condition decreases in BP and HR to reduce the pain of angina pectoris (*21*).

GARY E. SCHWARTZ
Department of Social Relations,
Harvard University,
Cambridge, Massachusetts 02138

References and Notes

1. For reviews of human research, see H. D. Kimmel, *Psychol. Bull.* **67**, 337 (1967); E. S. Katkin and E. N. Murray, *ibid.* **70**, 52 (1968). For reviews of animal research see N. E. Miller, *Science* **163**, 434 (1969); L. V. DiCara, *Sci. Amer.* **222**, 30 (January 1970).
2. D. Shapiro, B. Tursky, E. Gershon, M. Stern, *Science* **163**, 588 (1969); D. Shapiro, B. Tursky, G. E. Schwartz, *Circ. Res.* **26-27** (Suppl. 1), I-27 (1970).
3. J. Brener and R. Kleinman, *Nature* **226**, 1063 (1970).
4. D. Shapiro, B. Tursky, G. E. Schwartz, *Psychosom. Med.* **32**, 417 (1970).
5. A complete description of the ID model and the present experiment can be found in G. E. Schwartz, thesis, Harvard University (1971).
6. For this report, no formal distinction is made between feedback (implying response information) and reinforcement (implying response contingency). Since both terms require experimental procedures for systematically producing changes in the environment which closely follow changes in behavior, their similarities (in obtained results) rather than differences (in theoretical underpinnings) are emphasized. Although I have primarily used the terminology developed in operant conditioning, feedback terminology can be easily substituted; the major implications and conclusions remain the same.
7. These definitions depart somewhat from those generally used in biology, where the term integration refers to any consistent pattern of unified activity, regardless of direction, while differentiation refers to a separation of one response from others (here called specificity).
8. For example, if a response decreases naturally from the beginning to the end of an experiment this changing operant baseline must be taken into account to measure the direction and extent of learning. See A. Crider, G. E. Schwartz, S. R. Shnidman, *Psychol. Bull.* **71**, 455 (1969).
9. R. F. Rushmer, *Cardiovascular Dynamics* (Saunders, Philadelphia, 1961).
10. A similar approach has been used to train rats to differentially control blood volume between the two ears. See L. V. DiCara and N. E. Miller, *Science* **159**, 1485 (1968).
11. A preliminary experiment on learned cardiovascular integration in which the procedure was tested can be found in G. E. Schwartz, D. Shapiro, B. Tursky, *Psychosom. Med.* **33**, 57 (1971).
12. No subjects showing initial systolic blood pressures of 135 mm-Hg or more during the last adaptation or random reinforcement trial were included in this experiment.
13. Analyses of variance were performed on an IBM 360 computer with the Biomed 08V program. Groups was the between factor and trials was the within factor. The degrees of freedom for the group by trial interaction was 34/612. For the BP data, the analyses of variance (comparison of two groups at a time; six possible combinations) revealed that over trials the two BP^{up} conditions were each significantly higher than each of the two BP_{down} conditions ($P < .0001$), but were not different from each other. For HR, all two-group comparison trials were significant ($P < .01$ to .0001) except BP^{up} HR^{up} with BP^{up} HR_{down}, and BP^{up} HR_{down} with BP_{down} HR^{up}. The corresponding group main effects for these comparisons were also significant for BP ($P < .025$ to .005) and HR ($P < .10$ to .0001).
14. For example, separate analyses of variance on each of four groups, with measures (2) (change in BP versus change in HR) and trials (35) as within factors, indicated that BP_{down} HR^{up} reinforcement was the only condition that produced a reliable divergence between BP and HR during trials (d.f. = 34/306).
15. Sinus arrhythmia, a condition in which HR leads BP, is discussed by A. M. Scher, in *Physiology and Biophysics*, T. C. Ruch and H. D. Patton, Eds. (Saunders, Philadelphia, 1965), p. 660.
16. This pattern of results is also consistent with physiological theory suggesting that the parasympathetic system is capable of finer differentiation than is the sympathetic system. Raising the HR while lowering the BP may constitute a parasympathetic pattern, since HR may be increased by a decrease in vagal tone. Unlike elevation of HR, elevation of systolic BP requires sympathetic activity, hence the observed difficulty in lowering HR at the same time.
17. Finer analysis procedures (for example, coincidence measures of BP, HR, and respiration) may be necessary to assess such effects.
18. Predictive power of the ID model will, by definition, be limited to the extent that (i) operant (feedback) theory adequately handles the learning of individual responses (the interaction of other variables such as cognitive set and motivation is yet little understood); and (ii) physiological mechanisms and constraints can be empirically assessed in the given situation.
19. The recently published report by E. E. Fetz and D. V. Finocchio [*Science* **174**, 431 (1971)], which demonstrates operant conditioning of specific patterns of neural and muscular activity in the monkey, strongly supports this view.
20. For example, when changes at each beat in diastolic (as opposed to systolic) BP are reinforced, some conditioning of HR also takes place (D. Shapiro, G. E. Schwartz, B. Tursky, *Psychophysiology*, in press). This implies that diastolic BP and HR are partially (but not completely) integrated with respect to phase. Analysis of BP^{up} HR^{up} and BP_{down} HR_{down} coincidence responses has confirmed this prediction (average phasic integration is 70 percent). Research is required to determine the extent to which these two integrated functions can be separated through operant differentiation reinforcement (for example, requiring the subject to decrease diastolic BP by reducing peripheral resistance, while at the same time increasing HR).
21. The circumstances under which cardiac oxygen requirements in angina pectoris can be reduced if both HR and BP are lowered are discussed by E. Braunewald, S. E. Epstein, G. Glick, A. A. Wechsler, N. H. Wechsler, *N. Engl. J. Med.* **277**, 1278 (1967); E. H. Sonnenblick, J. Ross, E. Braunewald, *Amer. J. Cardiol.* **22**, 328 (1968).

22. Supported by NIMH research grant MH-08853; research scientist award K5-MH-20,476; ONR contract N00014-67-A-0298-0024; and the Milton Fund of Harvard. I especially thank D. Shapiro and B. Tursky for guidance and encouragement, and J. D. Higgins for comments on the manuscript.

4 April 1971; revised 8 July 1971

Backward Enhancement?

Abstract. *Presentation of a masking stimulus enhances, rather than detracts from, detectability of certain multisegment targets. Present theories of backward masking cannot account for this "backward enhancement" effect, which resembles another puzzling phenomenon, previously reported as target recovery or disinhibition. An explanation in terms of interaction between retinal excitatory and inhibitory fields is offered.*

Under proper circumstances the second of two successively presented visual stimuli can reduce, and sometimes entirely eliminate, the recognition and even the mere detection of the first (*1*). Depending on experimental paradigm, theoretical orientation, and personal preference, this retroactive inhibitory effect in perception is variously referred to as "metacontrast," "erasure," or "backward masking." For convenience, and following current practice, we use the terms, backward masking, or masking, in this paper. We refer to the first stimulus as the "target" and the second as the "mask."

Empirical investigations of backward masking have covered the effects of a host of temporal, luminance, and configurational variables. Several theoretical accounts of masking have been offered (*2*), some based on the neurophysiological concept of inhibition, and some more specifically on "lateral inhibition." The latter models have been quite successful in accommodating large portions of the reported empirical

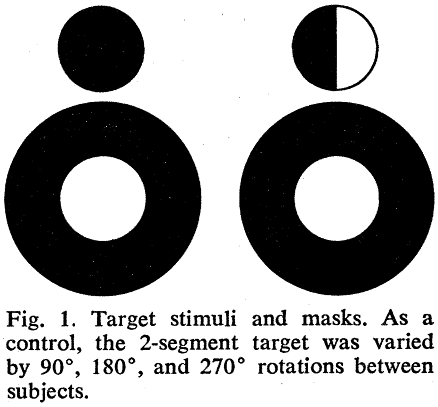

Fig. 1. Target stimuli and masks. As a control, the 2-segment target was varied by 90°, 180°, and 270° rotations between subjects.