DOCUMENT RESUME

ED 152 706 PS 009 843

AUTHOR Condry, John; Koslowski, Barbara
TITLE Can Education Be Made "Intrinsically Interesting" to Children?
INSTITUTION ERIC Clearinghouse on Early Childhood Education, Urbana, Ill.
SPONS AGENCY National Inst. of Education (DHEW), Washington, D.C.
PUB DATE 77
NOTE 48p.

EDRS PRICE MF-$0.83 HC-$2.06 Plus Postage.
DESCRIPTORS Early Childhood Education; Educational Strategies; Elementary Secondary Education; *Learning; *Learning Motivation; Learning Processes; Learning Theories; *Literature Reviews; Motivation Techniques; *Rewards; *Self Reward; *Student Motivation; Task Performance

ABSTRACT

This paper reviews the role of intrinsic and extrinsic motivation in children's learning. The paper contends that two types of learning exist: self-initiated learning which is intrinsically motivated, and learning initiated by another person which involves external rewards. The paper suggests that motivation is not unidimensional and that task performance reflects the motivational system involved. The paper reviews early and current research on the effects of rewards on activation which suggest that, for educational purposes, rewards limit a student's engagement in an activity, affect what is learned, and affect a student's desire to return to the activity involved. It is also suggested that rewards have undesirable effects on the teachers who dispense them. The paper contends that education should focus more on children's intrinsic motivation for learning. A variety of research is reviewed which highlights the child's intrinsic motivation to pursue questions of patterns and regularities in the world and related causal explanations. Also reviewed are studies which focus on the child's intrinsic desire to explore solutions to new problems and to relate new information to previously acquired knowledge. The paper suggests that education should take advantage of this intrinsic motivation to learn and to understand. (BD)

Can Education be Made "Intrinsically Interesting" to Children

John Condry and Barbara Koslowski

Cornell University

> All moral culture springs solely and immediately from the inner life of the soul, and can only be stimulated in human nature, and never produced by external and artificial contrivances.
>
> Whatever does not spring from a man's free choice, or is only the result of instruction and guidance, does not enter into his very being, but still remains alien to his true nature; he does not perform it with truly human energies, but merely with mechanical exactness.
>
> All peasants and craftsmen might be elevated into artists; that is, men who love their labor for its own sake, improve it by their own plastic genious, and inventive skill, and thereby cultivate their intellect, ennoble their character and exalt and refine their pleasures. And so humanity would be ennobled by the very things which now, though beautiful in themselves, so often serve to generate it. (Wilhelm von Humboldt, The limits of State action, in J. W. Burrow (Ed.), Cambridge studies in the history and theory of politics. Cambridge: Cambridge University Press, 1969, pp. 76, 63, 27-28. Quoted by Chomsky.)

Motivation has been a persistent problem in American education since compulsory education was adopted (Cremin, 1961). While there have always been children ready and eager to learn, there have also been those who were apparently disinterested, recalcitrant, and unwilling to "apply" themselves to their studies. Because our approach has been to study the organism and to speculate about its character, our theories have focused on how to "motivate"

the disinterested child. The assumptions of Behaviorist psychology were thought to be well suited to this task. The individual of any age, it was said, is motivated to attain rewards. Scheduled rewards contingent upon proper performance, were introduced into educational systems, and the result has been the token economy, and the current system of offering "grades" for proper school work.

While it is true that adults often work for rewards (a paycheck, for example) and it is not unreasonable to suggest that children may be motivated in the same way; this way of teaching has certain drawbacks, and we might be wise to study them in the hopes of discovering other processes that might be involved, and other means of educating children. A common disadvantage to extrinsic reward systems, for example, is the complaint that they motivate the child more to get the "reward" than to arrive at a complete understanding of the educational task (reading, math, science, etc.) at hand. In the

The material in this publication was prepared pursuant to a contract with the National Institute of Education, U.S. Department of Health, Education and Welfare. Contractors undertaking such projects under government sponsorship are encouraged to express freely their judgement in professional and technical matters. Prior to publication, the manuscript was submitted to the Area Committee for Early Childhood Education at the University of Illinois for critical review and determination of professional competence. This publication has met such standards. Points of view or opinions, however, do not necessarily represent the official view or opinions of either the Area Committee or the National Institute of Education.

1820's, for example, the society for Progressive Education in New York, introduced a system of redeemable tokens as rewards for correct school work and a system of fines for various offenses in the school. This was done in order to discourage corporal punishment, of which the society disapproved. This early version of the token economy was abandoned in the 1830's because the trustees of the society came to feel that "they were more often rewarding the cunning that the meritorious", and that the system of tokens "fostered a mercenary spirit." (Ravich, 1974)

In this chapter, we will contend and attempt to demonstrate that there are different learning processes which are dependent, in the first instance, on what or who initiates the process. There is one type of learning that is most common which we shall call "self-initiated" learning, and it is driven, we shall contend, by "intrinsic motivation". When rewards are offered in an attempt to motivate a child, a second type of learning is called into play, a type we shall call "other initiated" learning. A basic contention of this chapter is that the two forms of learning are different in their processes and their outcomes, and that one is a poor substitute for the other.

In the first part of this chapter, we will review evidence which suggests that the two processes just described are <u>not</u> part and parcel of the same motivational syndrome, but rather incompatable with one another. We will attempt to show that when rewards are offered, they change the nature and definition of the task, and eventually diminish the subject's "control" of the learning situation. A number of recent studies are reviewed and categorized all of which suggest that when performance on a task "for its own sake" is compared with performance initiated by the desire for an extrinsic reward, the subjects show different patterns of interaction with the task, and different motivational effects after the experience.

In the second part of this chapter we will attempt to develop a picture of intrinsic motivation in terms of the child's cognitive abilities or capacities for learning about the world. Much of the first part of this chapter is paraphrased from another article by the senior author (Condry, 1977, see also Condry, 1978, Condry & Chambers, 1978). We begin with a consideration of the background literature that led to the current round of research on the topic .. the effects of reward on motivation.

Background

The fact that "rewards" (grades, surveillance, exams) and other extrinsic incentives are disruptive of some kinds of motivation has been remarked upon for some time. Albert Einstein is quoted as having commented about exams: "This coercion had such a detering effect that, after I had passed the final examination, I found the consideration of any scientific problems distasteful to me for an entire year.". (Bernstein, 1973). A quarter of a century ago Harlow (1950, 1953; Harlow, Harlow, & Meyer, 1950) commented upon the inverse relationship between extrinsic incentives and the learning of complex problems in primates. Briefly, primates deprived of food (in order to "motivate" them) performed poorly on complex problems. Well fed (thus presumably "unmotivated") primates solved the same problems with ease. Shortly thereafter, White (1959), Berlyne (1955, 1957, 1958, 1966), Hunt (1965) and Koch (1956) all called attention to the distinction between extrinsic and intrinsic motives, and most (Koch is the exception) offered theories of intrinsic motivation. Basic to most of these theorists is the idea that intrinsic motivation is undermined by certain extrinsic conditions (see DeCharms, 1968). Thus, White (1959) suggested that anxiety was the "enemy of exploration," while Koch (1956) noted

that "any factor. . .that brackets the work on the task as instrumental. . ." will disrupt intrinsically motivated activities. These early theories suggested the existence of different motivational patterns which were, to some extent, defined by the circumstances of their initiation. Moreover, several theorists (especially Koch, White and DeCharms) felt that the different patterns of motivated activity were antagonistic.

The operation of extrinsic incentives was also called into question by Festinger and the dissonance theorists. In 1959, Festinger and Carlsmith published the findings of a study showing that subjects who were paid a small amount of money to lie about the interestingness of a dull experiment changed their attitudes more (in the direction of actually liking the experiment) than did a group of subjects who were paid a great deal more money. This negative relationship between incentive magnitude and attitude change has since been confirmed--subject to a variety of qualifications--in a number of studies (cf. Cohen, 1962; Lependorf, 1964; Carlsmith, Collins & Helmreich, 1966; Linder, Cooper & Jones, 1967; Helmreich & Collins, 1968; Nel, Helmreich & Aronson, 1965; and Bem & McConnell, 1970). Basically, the greater the extrinsic pressure exerted to "motivate" attitude change, the less real change is observed. Essentially the same effect for the interaction of powerful extrinsic incentives and internalization (this time for a prohibition) was found by Aronson and Carlsmith (1963), Freedman (1965), Pepitone, McCauley and Hamond (1967), and Zanna, Lepper and Abelson (1973). Children given a "mild" as opposed to a "severe" threat to motivate avoidance of a toy were less likely to play with the toy given a future opportunity to do so. Apparently there was greater internalization when the external pressure for compliance was less. All told, the line of research arising from dissonance theory raised another

series of questions about the interaction of extrinsic incentives and subsequent attitudes and behavior, and caused yet another crack in the facade of traditional motivation theory.

These two lines of evidence then: the one represented by Harlow, White, Hunt, Berlyne, and Koch and the other by the DeCharms, Festinger and the dissonance theorists, converge in the most recent research on this topic. The earlier theorists questioned the interaction of different forms of motivation, but the specification of the exact nature of this interaction and its ramifications had to await the current round of research and interest in the topic.

Current Research

The recent studies of the effects of rewards on motivation all have much the same general structure: first, a person's level of motivation for doing some task is either measured or assumed. This is the "base rate" measure of "intrinsic" motivation, usually defined as the person's willingness to do the task in the absence of task extrinsic pressure to do so (see Deci, 1975, Lepper, Greene & Nisbett, 1973). Next, some experimental intervention is attempted, usually involving a reward, and finally, a second attempt is made to assess a person's willingness to engage in the behavior under question. Any change in the level of interest from the first to the second measure is taken as the primary dependent measure of the effect of reward on motivation.

The first of these studies was done by Edward Deci (1971, 1972a) in an attempt to shed some light on the issues discussed above. Deci used a game called SOMA (trademark Parker Brothers) composed of a number of blocks which may be arranged to form a variety of patterns. The task in this research done with college students, was to reproduce certain patterns presented to the subject. During the middle of each of three experimental sessions, the

experimenter left the room and, unbeknownst to the subject, observed the subject's willingness to play with the puzzle in the absence of "demands" that he do so. In different experimental conditions, subjects were offered nothing, a monetary, or a social reward for every configuration produced with the puzzle. Deci's interest was in the amount of time spent playing with the puzzle during the "free" session, indicative of continued "intrinsic" interest in the task. In general, Deci found that (a) when money was used as an external reward, intrinsic motivation tended to decrease, and (b) when verbal reinforcement was used, intrinsic motivation tended to increase.

With interests similar to Deci before them (although they were apparently unaware of Deci's work at the time of their study), Kruglanski, Friedman, and Zeevi (1971) sought to clarify the same confusing literature on reward and motivation, and particularly the negative relation between the magnitude of incentive and subsequent liking for a task described by the dissonance theorists. Kruglanski et al. theorized that the "liking" found in earlier studies (e.g., Bem & McConnell, 1970; Carlsmith, Collins & Helmreich, 1966, etc.) is only one aspect of a more general syndrome of intrinsic as opposed to extrinsically motivated behavior. If so, these authors suggest, intrinsically motivated individuals: "might be expected to exhibit superiority on those aspects of performance contingent upon preoccupation with the task, as opposed to concentration upon attaining the goal." (emphasis mine, Kruglanski, et al. 1971, p. 607)

To test this proposition, 32 high school aged subjects were assigned at random to either a reward (called "extrinsic incentive") or a no reward condition. The rewarded subjects were offered a "guided tour of the Department of Psychology." The no reward subjects were offered nothing, and all subjects were given two measures of creativity, two measures of short term recall, and a measure of the recall of a series of incompleted tasks (the Zeigarnick effect).

On each of these dependent variables, subjects in the no reward group received higher scores than subjects in the reward group. The no reward group also liked the experiment slightly more and were slightly (but not significantly) more willing to participate in similar projects in the future.

The main finding of this study then, like Deci's, is that the offer of a reward prior to undertaking a task (or in this case a series of tasks) appears not only to lower subsequent interest in the task as Deci found, but also appears to lower the quality of performance on the task itself.

While both the Deci and Kruglanski et al. studies have powerful implications for education, both were done in laboratory settings. The next study to be reported, however, was done in a nursery school with "typical" incentives (for that situation). In this study by Lepper, Greene and Nisbett (1973), 51 nursery school aged children were selected based upon their demonstrated interest in a drawing task. Only high interest subjects were used, and these were exposed to one of three experimental conditions: In the expected award condition, subjects agreed to perform an activity in order to receive a reward (called a "good player" award comprised of a certificate and a gold star). In the unexpected award condition, the same reward was given to unsuspecting subjects after they completed the task, and in the no award group, subjects neither anticipated nor received an award. The critical measure of continued interest was the time spent playing with the drawing task during a subsequent "free play" observational period. Lepper, Greene and Nisbett found that the award undermined interest in the "anticipated" condition only, and that there was no significant difference between the unanticipated and the no award conditions in terms of later free play activity. This study not only introduces several methodological innovations of note, but it also raises some important questions for theoretical

consideration. It should be noted, for example, that the reward alone cannot be said to undermine anything, since the effect is obtained only in the anticipated reward condition, that is, only in the condition where the behavior is initiated by the offer of a reward.

From the three studies described so far, several findings emerge to challenge the traditional notions of the effect of reward on human motivation, and several issues are raised which require clarification. Apparently rewards do not always have salutory effects on motivation, the loss of interest in an apparently interesting and attractive task and the suggestion that the offer of a reward lowers quality of performance are serious and important matters.

In response to these initial studies and the issues they raised, a wide range of studies have sprung up to support and extend these early findings. In the sections to follow, we plan to categorize the most recent research on this topic, and begin an analysis of the issues involved.

In order to accomplish this, the remaining research will be described in terms of (a) the characteristics of the incentives studies, (b) the nature of the dependent variables, (c) the characteristics of the situations or tasks studied, and (d) the characteristics of the subjects studied.

Characteristics of the Incentives Studied

A. Expected vs. unexpected rewards. One of the most important facts to arise from this series of studies is that it is not the "reward" per se that affects subsequent interest, but the "timing" of the reward, in the sense of when it occurs. Rewards used to initiate task activity have different effects than rewards given after the activity is initiated by the subject. That is when the subject is aware of and anticipates a reward, his behavior is substantially different from that of comparable subjects who receive the same reward unexpectedly. This effect is clear in the research of Lepper, Greene and Nisbett (1973), and Smith (1974). In addition, Lepper and Greene (1975)

and Greene and Lepper (1974) have replicated these findings with different subjects a different task, obtaining the same results as in their earlier research. Deci's (1971, 1972a,b) and Kruglanski's et al. (1971) studies show substantially the same results when comparing rewarded to unrewarded subjects, but these authors do not use the unexpected reward condition, and so are not directly comparable in this respect.[2] The importance of this variable is that it focuses our attention on the context of the situation and not on the reward per se, and we shall return to this point later.

B. Contingent vs. non contingent rewards. The contingency between the reward and performance on the task is another issue that arises in the current research. Some of the early work has made the reward contingent only on "doing the task" (Lepper et al. 1973; Ross 1975; Upton, 1973; Lepper & Greene, 1975; Kruglanski et al. 1971; Deci, 1971; Benware & Deci, 1974), while other studies have drawn a more explicit contingency between the quality of performance and the reward (Deci & Cascio, 1972; Deci, 1972a; Smith, 1974; Garbarino, 1975; McGraw & McCullers, 1974; Ross, Karniol & Rothstein, Note 1; Karniol & Ross, Note 2). Early research on this topic reached apparently opposite conclusions about the effects of contingency on performance. Deci (1972b) found no undermining effect for noncontingent rewards, but rewards contingent upon performance did produce an undermining effect similar to that found in most other research. Greene and Lepper (1974) on the other hand, employed contingent and non contingent rewards in a replication of the earlier Lepper et al. (1973) study, and found both contingent and non contingent rewards had the effect of undermining future interest. One difficulty in interpreting findings across a number of different studies is that researchers often use the same word to describe different events. Thus the "non-contingency" employed

by Deci (1972b) refers to rewards which were unrelated to the task (e.g., paying subject at the outset for participating in the experiment), while the non-contingency in the Greene and Lepper (1974) study refers to rewards for "doing the task" but not explicitly tied to a performance criterion.

One interesting possibility that may account for the inconsistency in the early research on this topic could be that most of the researchers who have used contingent rewards have done so with adults, while most of the studies using a non contingent design have done so with children. It may be difficult for a young child to completely understand contingency instructions ("We do not have enough rewards to give everybody one, so we will only give awards for very good drawings. . ."). First, nursery school children are seldom put in the position of having their work judged and "rewarded" on a contingency basis. A child who has experienced a relatively undifferentiated environment, especially with respect to art work, might not be able to judge his work comparatively and thus be unresponsive to contingency demands. Secondly, many young children have not yet developed the "adult" usage of comparatives (witness the child who gets into a hot bathtub and says. . ."It's too hot, make it warmer".). Yet contingency manipulations depend, for their successful employment, upon understanding comparatives. Finally, in addition to being unable to discriminate between "good" and "very good" work, young children may not possess the conscious control over their behavior to be able to produce a "very good drawing" when called upon to do so. When children are told to look at something "in order to remember it," they don't do anything different than when just told to look at it. They don't remember it any better either. Adults do use different strategies in these same circumstances, and they do remember better when told to do so (Reisser, 1976).

Skills and strategies which have not yet developed and which are not yet stable cannot be "produced" on demand. The most recent studies on this topic Ross (1975) have traced the decrement in interest to contingent reward situations.

C. Salience of reward. In two studies the "salience" of the reward was manipulated, and the effects of this procedure examined. Ross (1975) manipulated the salience of the reward directly by leaving the anticipated reward in front of the subject (albeit under a box) while the task was being done. Undermining effects were observed for the salient reward condition only. Although the same reward was offered and given to the non salient group, it was simply not available during the task. In a similar design with a different purpose, McGraw and McCullers (1974) gave children rewards (M & M's) in two conditions, one where they were given a reward for each correct response (it was dropped in a bowl beside the subject) and in the other condition the same reward was used to "mark" a correct response, but the child was told he would not be allowed to consume the candy (it was simply a "token" of correctness).

Subjects in these two conditions did more poorly in a discrimination learning task than children in a third (control) group given neither rewards nor "tokens." Subsequent interest was not measured in this study, but the demonstration is important for the suggestion that the effect of reward was to undermine actual performance on a task, and the more salient the reward, the more undermining of performance was observed.

D. Type of reward. The rewards used in the research described herein are many and varied among studies, although there have been few attempts to vary the nature of the reward within a given study. Monetary rewards are common with adults (Benware & Deci, 1974; Deci, 1974; Deci, 1971, 1972a,b;

Smith, 1974; Upton, 1973), while tokens, certificates, and edibles of one sort or another are more common with children (Garbarino, 1975; Greene, 1974; Greene & Lepper, 1974; Kruglanski, Freedman & Zeevi, 1971; Lepper & Greene, 1975; McGraw & McCullers, 1974; Lepper, Greene & Nisbett, 1973). We have already seen how more salient rewards produce stronger effects than less salient ones, so it stands to reason that the stronger the incentive, the more powerful the undermining effect, although no manipulations of this sort have been tried using a paradigm that looks at continued interest as the primary dependent variable.

1. Monetary vs. social rewards. One interesting conflict to emerge from this research concerns the role of social rewards, or, as Deci calls it, feedback, as compared to monetary rewards. In a series of studies (Deci, 1971, 1972a, Deci, Cascjo & Krusell, 1975) Deci has shown that positive feedback (praise) can enhance subsequent interest for males, while it seems to have the opposite effect for females. Smith (1974) was unable to replicate these findings when he used a "Lepper type" (Expected, Unexpected & No Reward) design for both social and monetary reward. Smith found that both types anticipated of reward undermined future interest for both sexes. It is difficult if not impossible to resolve this issue with the evidence at hand. Deci used unanticipated social reward only whereas Smith used both anticipated and unanticipated social reward. In addition, the tasks used were different. The task used by Smith (learning about art) might have been seen as less sex appropriate than the SOMA puzzle used by Deci. There is evidence from research with children that they both see tasks and activities as sex typed (Hartley & Hardesty, 1964; Stein & Smithells, 1969), and evidence to indicate that information about the sex appropriateness of activities affects children's

responses to them (Liebert, McCall & Hanratty, 1971, Montimayor, 1974; Stein, Polky & Mueller, 1971).

2. Surveillance. Two studies are worth mentioning for the light they shed on other contextual manipulations that appear to have the same undermining effect as the "positive" incentives used in most studies. Lepper and Greene (1975) in a clever variant, used a condition of "surveillance" and were able to produce undermining of interest effects similar to those found with the various reward conditions discussed above. In this study, nursery school children did a task while being observed by a TV camera "on" more frequently in high than in low surveillance conditions. Those in the surveillance conditions were less interested in the task in a later free choice period than children in a no surveillance condition.

In a study done in the context of a different line of research, Zivin (1974) had adults try to interest distractable children in a boring toy. Contrary to her expectations, the treatment had no such effect. Children told to "think about interesting things you could do with this toy" played no more with the target toy than children who were given no adult encouragement.

These findings are particularly interesting when one takes into consideration Zajonc's (1965) review of social facilitation. Zajonc reviewed 50 years of research on the effect of the presence of others on learning and performance and came to the conclusion that the simple presence of other people tends to have the effect of undermining the development of a poorly learned skill, while it facilitates the performance of well learned skills. Based on this observation, then, we are led to suggest that like expected reward, the presence of others in the surveillance conditions changes the context of the learning situation in ways that are detrimental to task performance as well as subsequent interest.

3. Negative incentives. Finally, while most of the research has focused upon "positive incentives," at least two researchers have analyzed negative feedback and threats. Deci and Cascio (1972) looked at changes in intrinsic motivation as a consequence of threats and negative feedback (failure). These researchers found that "negative" incentives also decrease intrinsic motivation. Weiner and Mander (1977) studied the effects of either a shock or the anticipation of a shock on subsequent interest. Subjects who were shocked at random while performing an anagrams task showed greater willingness to pursue the task (when neither the shock nor the experimenter were present) than subjects who received no shock during the manipulation phase of the experiment.

Earlier research on the severity of initiation done by dissonance theorists (e.g., Aronson & Mills, 1959; Gerard & Mathewson, 1966) found the more severe (negative) the initiation, the greater the subsequent attraction to a group on the part of subjects.

The implication of these studies is that even negative incentives, when used as sanctions against unwanted behavior, have complex effects. If a negative sanction (threat or punishment) is not effective in stopping the behavior, may actually increase a subject's interest in and attraction to the prohibited behavior.

In sum, a variety of anticipated, salient, positive incentives (the primary independent variable in this research) and some additional task extrinsic conditions (surveillance and negative incentives) have been shown to be associated with lower subsequent interest in a task when compared to conditions where no task extrinsic incentives are available. The fact that the same incentives given after the performance of the task (and unanticipated by the subjects) produce no such "undermining" effects suggests that it is

the "context" created by the offer of rewards and not the rewards themselves which is responsible for these findings.

These facts alone are contrary to theories of motivation which assume extrinsic & intrinsic motivation to be additive, as both Deci (1975a) and Calder and Staw (1975) have noted. In consequence, a number of researchers have offered theories to account for the decrement of interest findings (cf. Kruglanski, 1975; Deci, 1975; Green & Lepper, 1975; Calder & Staw, 1975), and some have denied the validity of them (Reiss & Sunshinsky, 1975; Scott, 1976). But before it is useful to theorize about these effects (or to deny their existence) the entire range of effects should be outlined and described. In addition to the undermining of interest we have been describing, there is an additional suggestion from some of this research that rewards may have detrimental effects on both the process of learning and the products of learning, and for a consideration of these effects, we now turn to an analysis of dependent variables studied, and the tasks or situations utilized in the current research.

Nature of the Dependent Variables

A. Subsequent Interest. In most of the foregoing research, the dependent variable of primary importance was the degree of subsequent interest in the task shown by the subjects. Two different measures of subsequent interest have been used: attitudinal measures of liking for the task and/or willingness to return at a later date (e.g., Calder & Staw, 1975; Kruglanski, Alon & Lewis, 1972; Kruglanski, Ritter, Amitai, Margolin, Shabtai & Zaksh, 1975; Kruglanski, Ritter, Arazi, Agassi, Monteqio, Peri & Peri & Peretz, 1975) and behavioral measures of persistence in the activity in the absence of "extrinsic" demands (e.g., Deci, 1971; 1972a Lepper, Greene & Nisbett, 1973; Green & Lepper, 1974; Leeper & Greene, 1975). In some cases, both attitude and behavior measures have been used (Smith, 1974; Ross, 1975; Karniol & Ross,

Note 2; Greene, Sternberg & Lepper, 1976). In general these measures show similar effects and the choice of what type of dependent measure to use is parimarily of theoretical interest.

B. Quality of performance. In addition to measures of subsequent interest, some measures of quality of performance during the task have been taken, and the results suggest that anticipated rewards lower the quality of activity during a task as well as the desire to return to a task at a later time. Kruglanski, Freedman and Zeevi (1971), for example, found the promise of a reward effected "qualitative" differences in task performance, for memory, creativity and the recall of incompleted tasks (The Ziegarnic effect), and Lepper, Greene and Nisbett (1973) reported lower quality drawings for the "expected" reward group. Greene and Lepper (1974) replicated this result for the quality of drawings in the anticipated reward condition in a second study, and they found, in addition, that the expected reward group produced quantitatively more drawings. These findings suggest that one of the effects of anticipated rewards on task performance is to increase activity but to lower the quality of that activity when compared to subjects doing the same task for no reward.

McGraw and McCullers (1975) report a number of studies with children which find that tangible rewards, given on a trial-by-trial basis, lead to more errors and less learning when the performance of rewarded subjects is compared to nonrewarded subjects. These tasks range from perceptual discrimination (Miller & Estes, 1961), concept identification (Masters & Mokros, 1973; McCullers & Martin, 1971; Terrell, Durkin & Wiesley, 1959) verbal discriminations (Haddad, McCullers & Moran, noted in McGraw & McCullers, 1975, p. 14; Spence, 1970), picture discriminations (Scherc 1969; Spence & Dunton, 1967; Spence & Segner, 1967), and a patterned probability task (McGraw & McCullers,

1974). In short, in the words of McGraw and McCullers (1975); "within the confines of children's discrimination learning. . .the detrimental effece of reward is a very general one." (p.3)

In addition to discrimination learning, McGraw and McCullers have been able to extend their analysis of these effects on performance to probability learning (McGraw & McCullers, 1974; McGraw & McCullers, Note 4) and incidental learning (Staat & McCullers, Note 5). Taken together, these studies suggest that task performance is disrupted more for subjects who are given task irrelevant rewards than for comparable subjects performing these same tasks for no reward. Subjects given tangible incentives make more errors, solve the problems more slowly, make more stereotypical responses and do more poorly in S-R recall (in an incidental learning task) than do subjects who are offered or given no reward.

Characteristics of the Tasks or Situations Studied

Few of the studies reported have discussed the ongoing process of task performance. In part, this is due to the fact that the tasks used by researchers have not been ones where characteristics of process can be easily assessed. Examples of actual tasks used in the research have ranged from a puzzle called SOMA (Deci, 1971, 1972a,b; Deci, Cascio & Krusell, 1975), to drawing with a felt tipped pen (Lepper, Greene & Nisbett, 1973; Greene & Lepper, 1974), from solving a simple concept formation problem (Lepper & Greene, 1975; Garbarino, 1975; Maehr & Stallings, 1972) to pro-attitudinal advocacy (Benware & Deci, 1975), and last, and probably least, beating a drum (Ross, 1975).

A. Interesting vs. dull and boring tasks. While much of the recent research has utilized tasks chosen to be interesting (and thus "intrinsically

motivating") it will be recalled that earlier research by dissonance theorists tended to use dull and boring tasks in order to make theoretically similar points. Thus, as mentioned earlier, the dissonance theorists were able to show how a dull task could be made "more interesting" by inducing the subjects to lie about it for a small (as opposed to a large) reward. Contrawise, the intrinsic motivation researcher has shown how an interesting task may be made boring, with the introduction of a reward.

Few of these studies have varied the nature of the task within the experiment. One study to do so finds results which appear to be at variance with the earlier dissonance research. Thus, Calder and Staw (1975) find that while rewards decrease interest in "interesting" tasks, they may "enhance interest in dull and boring tasks." More research is needed to clarify the circumstances under which a dull task may be made more interesting.

B. The process of learning. In part as a consequence of the tasks used, most of the foregoing research is couched in terms of the products of learning, i.e., the number of problems solved, the number of errors made, the number of items recalled, etc. But it is also possible to study the process of learning in terms of the step-by-step strategies of action that subjects take. When this is done, even more of the picture emerges. Condry and Chambers (Note 6), for example, using a problem solving task originally designed to study strategies of thinking (Bruner, Goodnow & Austin, 1956), found that rewarded subjects attempted easier problems, required more information before they achieved a correct solution (i.e., were inefficient), were more "answer" oriented and less logical in their problem solving strategies than comparable nonrewarded subjects. Similarly, Maehr and Stallings (1972) found that children who believed they would be "evaluated" by an "external" source chose easier problems than children who believed they would evaluate their own problems

("internal evaluation"). Thus, there is evidence that extrinsic incentive conditions lead subjects to different strategic activities in a learning or problem solving situation than condictions which encourage exploration without the offer of a task extrinsic incentive as a reward.

Apparently these same "process" and strategy effects extend to interpersonal interaction as well as a person's interaction with a task. Garbarino (1975) looked at the effects of anticipated reward on children teaching other children, in a field setting experiment. Based on some of the "early" research reviewed above, he predicted that if a child were offered a reward to teach another child something, the quality of the interaction would be more "obtrusive, concrete, direct, and impatient in response to frustration" than in a situation where one child reaches another where no offer of a reward is made. In essence, Garbarino argued, like Kruglanski et al. (1971) before him, that the child in the "reward" condition would adopt an "instrumental" approach toward the situation, and focus more on the "goal" of receiving the money and less on the task of teaching itself. Specifically, Garbarino hypothesized that the tutor in the no reward condition would be more positive in her response to her tutee (all subjects were female), more efficient and less intrusive in her teaching style. He predicted that the younger child (the tutee was a younger child in all conditions) would learn more when her tutor was in the no reward condition.

The subjects in this experiment were 48 female elementary school children. Twenty-four 5th and 6th graders acted as tutors, for 24 1st and 2nd graders. The school in which the experiment was done was one in which cross age tutoring is an everyday fact of life, so the experimental situation was not an unusual one.

The older children were taught a game (a simple symbol substitution task) by the experimenter and then asked to teach it to a younger child, under either reward or no reward conditions. While the older child taught the game to the younger child, the interaction was observed and coded by the two experimenters. The results of Garbarino's study were supportive of his predictions and they add an important dimension to the work reviewed herein. The tutor's behavior toward the young child was more "negative" in the reward condition, and more positive in the no reward condition. In addition, the tutors in the reward condition made more demands upon the younger child than did the tutors in the no reward condition.

The promise of a reward not only affected the people to whom it was offered, but it also had secondary effects on the learning of the younger children (who were offered nothing). The younger children taught by tutors in the reward conditions scored lower for their task ability and higher for number of errors than did children taught by tutors in the no reward condition.

Garbarino's research substantially extends the range of the effects we have been observing. His findings suggest that in addition to undermining continued interest, the process of exploration and the products of learning promise of a reward has an important effect on the context of an interdependent, interpersonal situation. Apparently when the actions of one person may interfere with an extrinsic goal of another, a negative evaluation of the task is translated to a negative evaluation of the person. The instrumental orientation among tutors in the reward condition suggests that they valued the younger child as a "function of her utility in obtaining the desired goal and devalued her in proportion to the degree to which she failed and thus frustrated the tutor." (Garbarino, 1973)

C. Reward endogenous vs. exogenous tasks. A final point regarding the nature of tasks used in this research is worth mentioning. Kruglanski has shown in a number of studies (cf. Kruglanski, Riter, Amitai, et al., 1975; Kruglanski, Riter, Arazi, Agassi, et al., 1975) that when money is endogenous to a task (e.g., a game like coin tossing where, traditionally, the winner keeps the money; or poker) its presence "enhanced" intrinsic motivation, whereas when it was exogenous to a task (e.g., doing a jig-saw puzzle), its presence "depressed" intrinsic motivation. Essentially, the same conclusion is reached by Staw, Calder and Hess (Note 8). Most of the incentives used in the research described above were ones not commonly associated with the tasks, but the point is well taken from a theoretical point of view. For this reason, throughout this paper I have referred to the effects of task-irrelevant rewards and incentives.

Characteristics of the Subjects

A final area of concern for the current research is variables associated with the subject. One would think this is the first place to look for differences in the effects of rewards and other extrinsic demands, since neither the rewards themselves nor the tasks exist independently of the value the subject places on them (in terms of rewards) or the interaction he has with them (in terms of tasks). Yet this important area for research has received relatively little attention.

A. Personal characteristics of subjects: age, sex, and personality. The effect of rewards in undermining interest has been demonstrated on a wide range of ages. The studies reviewed have demonstrated the effect on nursery school children (Greene & Lepper, 1974; Lepper, Greene & Nisbett, 1973; McGraw & McCullers, 1974; Ross, 1975; Zivin, 1974), elementary school children (Garbarino, 1975; Greene, 1974; Kruglanski, Alon & Lewis, 1972; Maehr &

Stallings, 1972; McGraw & McCullers, 1974; Zivin, 1974), high school students (Kruglanski, Freedman & Zeevi, 1971), college students (Benware & Deci, 1974; Deci, 1971, 1972a,b; Deci, Benware & Landry, 1974; Deci, Cascio & Krusell, 1975; Smith, 1974) and adults beyond college age (Kruglanski & Cohen, 1973; Upton, 1973).

Aside from demonstrating the effect on a large range of subjects, however, few studies have attempted to explore other subject variables that may interact with the reward context and/or with the task. Subjects of both sexes have been used in many of the studies reported, but few significant differences due to sex are reported, with the exception of Deci's (1972a; Deci, Cascio & Krusell, 1975) finding that females react differently to positive verbal feedback, and Smith's (1974) findings to the contrary. These results were discussed earlier in terms of the interaction of sex with the reward and task characteristics.

Two investigators have looked into the effects of personality. Maehr and Stallings (1972) found that subjects high in Need Achievement volunteered more for difficult tasks when the evaluation context was internal, and more for easy tasks when the evaluation context was external. In addition, Haywood and his colleagues (Haywood, 1971; Haywood & Switzky, 1971) have developed a personality test for intrinsically motivated (IM) as opposed to extrinsically motivated (EM) individuals. Switzky and Haywood find that personality interacts with the reinforcement context. In this study, two reward contexts were studied: Self-reward, where performance standards are set and rewards delivered by the subjects themselves, versus external reward where standards are imposed and reinforcers externally administered. Under these circumstances, IM children "maintained their

performance longer than EM children under self-reinforcement, while EM children showed greater performance maintenance than IM children under external reinforcement". Clearly individual differences exist which modify the effect of each of these reward contexts, and these need to be studied.

B. Initial interest. One of the most fascinating, yet least studied, factors concerns the effect of reward upon people who vary in initial interest (Condry, 1975). Attractiveness of the task was considered earlier, in this section we consider initial interest of the subject, with the task held constant. Yet most of the researchers to date have selected for study only those who are high on interest. This leaves unanswered the question of whether rewards may have different effects on people with low initial interest. In all of the studies reported, only three have addressed this topic.

Lepper, Greene and Nisbett (1973) did a re-analysis of their initially high interest sample by dividing it at the median on interest. They found an increase in subsequent interest among the lo interest subjects only in the unexpected reward condition. Greene (1974) also studied subjects who varied in initial interest. Using a "between" treatment groups analysis, Greene (1974) found an overjustification effect for low interest subjects but not for high interest subjects when these groups were compared to a control. Using a within group analysis, he found a significant post-treatment decrement for high interest subjects, but not low interest subjects. Thus, even though Greene attempted to study this issue, this particular comparison was ambiguous because the data from the control subjects did not remain stable over time, and the between and within subject analyses did not agree (Greene, Sternberg & Lepper, 1976).

Finally, Upton (1973) studied the offer of a monetary reward as it effects people's willingness to donate blood to a blood bank. Subjects were divided on interest in terms of the number of pints of blood they had donated in the last 12 months. At this point, half of the subjects in the low interest group and half of those in the high interest group were offered and later given a reward of $10.00 for donation of a pint of blood. The other half of the sample was asked to donate, but no money was offered. The results indicate that subjects high in initial interest are significantly more willing to donate blood when they are not offered a sizable incentive to do so. Subjects in the low interest group go slightly, but not significantly, in the other direction, that is, they are equally willing to donate when money is offered or not. So, we have some evidence that unexpected rewards may increase interest in low interest subjects, but on the whole, the question of the effects of reward on low interest subjects is still very much up in the air.

All in all, the evidence described above suggests that task extrinsic rewards, when they are used to "motivate" activity, particularly learning, have widespread and possible undesirable effects. These extend to effects on the process as well as the products of the task activity, and to the willingness of the subject to undertake the task at a later date. It is difficult to summarize this material adequately, but in general, compared to nonrewarded subjects, subjects offered a task extrinsic incentive choose easier tasks, are less efficient in using the information available to solve novel problems, tend to be answer oriented and more illogical in their problem solving strategies. They seem to work harder, produce more activity, but the activity is of a lower quality, contains more errors, is more

stereotyped and less creative than the work of comparable nonrewarded subjects working on the same problems. Finally, to return to the point of departure, subjects are less likely to come back to a task they at one time considered interesting after being rewarded to do it. The facts appear true of a wide range of subjects doing a wide range of tasks. Attempting to account for them is a formidable challenge.

What these data suggest is that there are different forms and patterns of motivated activity. Learning under instruction while being guided, supervised, and directed from without is obviously of value. Too many highly skilled activities require it to dismiss this form of learning. But it is not the only form of learning we have observed, and if subsequent interest is a goal of learning, it is not even the most efficient.

The studies just reviewed suggest that the conditions of initiation are immensely important, in terms of both what the individual puts into a task, and what is gotten out of it. When a person chooses to engage a task, his behavior during the task is more coherent and his subsequent interest in the task remains higher than a comparable individual who is "pressured" into doing the same task by the offer of task extrinsic rewards.

It is easier to understand the meaning of all this research if we imagine the ecological circumstance we are trying to understand. Perhaps the entire pattern of these results would be seen more easily if we imagine motivated activity as having at least four discriminable states or phases as follows: a.) initial engagement; b.) activity or manipulation (the process); c.) disengagement; and d.) subsequent engagement. That is, it is possible to ask about the "forces" that act in each of these phases and of the relations between actions in one phase and another. We would want to know why a person engaged a task in the first place, what is done with it or to it while one is actively "manipulating" it, what leads one to disengage

it a ' possibly return to it. Such an analysis is useful since it fits the example of skills learned and utilized in everyday life and so gives us a reference point from which to judge the literature just described.

From this perspective, we may say that most of the research reviewed herein has focused upon the relationship between the conditions of task engagement (Phase I) and the degree of subsequent interest (Phase IV). Some research has looked at the relation between Phases I and II, that is, how the conditions of engagement affect the manner of activity and manipulation (including the manipulation of other people as in Garbarino's study). Few studies of disengagement (Phase III) are found in the current literature, but in the original research on the Ziegarnik effect, Ovsiankina (1928) studied what happens when one interrupts activity before it reaches its natural end and found a strong desire to return to the interrupted behavior (see Ryan, 1970, p. 96ff). In most of the research described earlier, the extrinsic conditions are "ended" when the reward is given. In the "intrinsic" conditions, however, the task actively is also interrupted and ended by the experimenter, i.e., before the individual "chooses" to leave it. It is not clear, from the research presented here, whether a self-generated disengagement might produce a different pattern of subsequent interest. An analysis by the various "phases" in the learning process allows us to see gaps in the research as well as organize a good deal of what has been done.

Putting these findings in perspective we may arrive at a different ew of the child in the educational system. The creature we hope to educate is complex, coherent and well prepared to learn about the world. Moreover, the creature is eager to learn, thirsty for knowledge. The danger signaled by these findings is that by using "smiling faces", stars, and eventually grades in order to motivate children to learn, we unintentionally limit the utility of what is learned, and undermine the child's natural curiosity.

What are the implications of a different view? Instead of asking how one might manipulate a disinterested child so as to "motivate" him to learn, why not ask how we may arrange the environment of education so as to take advantage of the child's natural curiosity and his intrinsic interest in learning about the world? Before children enter school they acquire a vast range of knowledge about the world. They were not "trained" to acquire this information in the sense of being supervised, scheduled, rewarded and punished, but rather they used their native, intellective capacities. This interest did not have to be "encouraged" for the most part, it was there, within, all along.

What are the elements of this intrinsic motivation? How does it function? In the next section of this chapter, we address this question in terms of the child's capacities for learning. It is clear that other people can do much to enhance the child's knowledge of the world by virtue of an awareness of the variety of conceptual tools the child has for understanding the world. We will not be describing specific curricula or concrete activities. These are best done by the people who are actually working in the classroom. What we will be describing are the facts of research findings that might be profitably extended to actual classroom situations. At present, we have no direct evidence that such extensions would be either practical or effective. However, lest we paint too bleak a picture, it is worth pointing out that at least some of the observations reported are of children engaging in activities and working with materials which, though not actually found in the classroom, at least bear a strong resemblance to the sorts of activities and materials that might occur in classroom situations.

Intrinsic Interests and Capacities

The argument that extrinsic rewards are of limited utility in educational settings is a useful argument only if it can be shown that the child himself is both intrinsically motivated to learn on his own and has the basic capacities to carry out what he is intrinsically motivated to do. Without intrinsic motivation accompanied by the appropriate capacities, a reliance on extrinsic rewards may be the only option open to an educational system.

In order to argue that children are intrinsically motivated to learn about important aspects of the world and are equipped with a capacity to do so, we will focus on findings from two areas of research: The first deals with the way in which children make sense out of or explain their environments; the second, with the process that is involved in solving nonroutine problems. The primary reason for this choice is that issues examined in this research are not bound to learning in particular content areas. They cut across and underlie learning about arithmetic and reading as well as about social studies and history. They occur when the child is engaged in subjects that are part of a formal curriculum as well as when he is engaged in trying to make sense out of the various social rules and conventions that govern his behavior with other people. These areas were chosen for another reason, as well. Implicit in the first part of this paper was the premise that there is more to learning than acquiring the skills necessary to churn out the correct answer. Learning is much deeper and, in the long run, substantially more useful, when one also understands the underpinnings of the correct answer, the process which enables one to understand why the

correct answer is correct. One can correctly solve (i.e., correctly apply the mechanical formula for solving) quadratic equations or correctly notice that certain events routinely co-occur without having a very thorough understanding of what quadiatic equations are all about or of why various events frequently co-cccur. The research areas we have chosen speak to the question of how the child learns about process.

One of the most unequivocal activities that children are intrinsically motivated to engage in is the enterprise that consists of organizing and explaining the world. In its most basic form, this enterprise involes the detection of organization and regularity that actually does occur. In a more sophisticated form, it involves imposing organization on, or inferring it in, situations in which it either does not exist or else exists in a less than perfect form. But the child does not limit himself to attempts to organize the world; he asks in addition why various regularities exist as well as how a particular kind of regularity or organization comes to be. In short, he searches for explanations for, as well as instances of, organization in the world. And, although his explanations may be false, they do nevertheless follow certain rules.

The Concern with Environmental Regularity

Evidence of the child's concern with environmental regularity comes from a wide variety of behaviors and situations. Some of the evidence was provided by children's spontaneous behavior, some by behavior that was elicited. For example, researchers such as Watson (1977) and Bower 1974) (to name only two) have found that even very young infants will quickly

learn to detect the regularity with which one of their own responses (such as sucking, head turning or limb movement) is followed by a particular effect (such as the movement of a mobile or the appearance of a person who says, "Peek-a-boo.") Infants will also detect regularities between two types of events when neither type consists of one of the child's own actions, i.e., when both events are external to himself. For example, if an object goes behind a screen and reappears on the other side, the infant will soon come to look at the other side of the screen before the object reappears there (Bower, 1974). By anticipating the reappearance, he gives evidence of having detected the regularity. Similarly, if an object appears first in one window and then another, an infant will soon anticipate its appearance in the second window (Anglin and Mundy-Castle, cited in Bower, 1974). The fact that even young infants detect environmental regularity is some evidence of how basic this tendency is.

Findings from studies of older, preschool children show an adeptness at dealing with even more subtle instances of regularity. For example, children of this age can easily perform concept attainment tasks that require them to detect the systematic occurrence of a common element in a number of different contexts (Vinacke, 1952). Indeed, children can detect environmental regularity even when the regularity is not perfect. For example, children are able to learn the correct response in probability tasks, tasks in which reinforcement follows the correct response only some of the time (Stevenson, 1970).

In addition, there is evidence that, once children do detect some modicum of organization or regularity, they will spontaneously extend what they have detected. In doing so, they often structure aspects of the world in a way that is actually more organized than the structure that in

fact exists. For example, Ginsburg (1977) describes a child who is in the process of learning to count and who extends the rule for forming tens numbers in the following way: "Ten, twenty...ninety, tenny." In the same vein, there is the, by now, classic example of children who have just learned to form the past tense by adding the suffix ed to verb stems. Such children will frequently also produce such constructions as "ranned," in spite of being exposed only to "ran." Finally, in some cases, children will actually invent their own way of organizing the environment. Ginsburg (1977), for example, reports that a not uncommon belief among grade school children is that, in adding columns of two-digit numbers, when the sum of the first column is itself a two-digit number, one carries over to the next column the larger number of the two-digit sum (i.e., the 9 in 19) rather than the number that is on the left.

At this point, we cannot resist reporting an anecodote that illustrates an additional way in which the tendency to impose organization on the environment manifests itself in the real world. A fourth grader came to the conclusion that the names of national leaders reflected the names of the countries which they represented. Her evidence was that: Kenyatta was the head of Kenya; deGaulle was the head of France (which, after all, used to be called Gaul); and Franco was the head of Spain. (Having heard that the early inhabitants of Europe were often nomadic, he assumed that Spain had at one time been inhabited by Franks.) Of course, this theory did not hold for the United States, but then maybe that's why they called the United States a melting pot; things got mixed-up in it. But there were vestiges of this correlation even here: Witness, e.g., George Washington and Washington, D.C. Now, what this theory lacked in accuracy (or even coherence) it made up for in the extent to which it imposed organization on what, in the absence of this theory, was nothing more than a collection of facts.

The child's acquisition of knowledge about the social world follows a similar pattern. The child begins to learn about gender identity for example, by first categorizing the world into two clusters (male & female) and then arriving at stereotypical views of the roles occupied by each sex, (Condry 1978, Kohlberg, 1975). This device allows for the child a form of control by knowing -- in a broad way -- what to expect. Control involves both the ability to anticipate the correlational structure of the world, and to manipulate the causal structure of the world.

There is fairly extensive evidence, then, that one way in which children try to achieve control over their worlds is by organizing them. If the organization is built into the world, children will detect it. If the organization is less than perfect, they will improve upon it. If the intended organization is not detected, they will construct and impose their own version. Children come equipped with a tendency to gain control over a large number of discrete (and possibly even unrelated) pieces of information by reducing them to particular instances of rules that are more general, and thus fewer in number, than the individual facts that they subsume.

Intrinsic Interest in Causal Explanations

However, the intrinsic motivation, if you will, to detect regularity in (or impose it on) the world is frequently not enough to satisfy a child. Often, he seeks to know, as well, why and how this regularity is brought about. Detecting regularity provides the child only with the information that various events in or aspects of the environment are associated with one another. It provides him with information only about correlation. It does not provide the child with any information about why events are associated, about how it is, for example, that one event regularly co-occurs with or follows another. Answering these sorts of questions requires causal explanations for the way the world is organized. The child must

move from correlation to cause. Sometimes explanations for why events are associated consist of no more than an identification of exactly which aspect of the environment is functioning as the causal agent in a particular situation. For example, the child may note that the wheels of a bicycle turn is because the pedals are pushed; the reason that lights go on is because switches are turned. This information in itself makes the environment more manageable because it enables the child to learn that particular phenomena are associated with only some aspects of the environment rather than others. Often, however, the child also becomes involved with the underlying mechanism or causal underpinnings that explain how it is that a particular causal agent (such as a switch) is able to make an event happen. That is, he concerns himself with questions about the intervening connection that mediates between causal agent and the effect and, thus, enables the causal agent to bring about the effect. For instance, in the example just mentioned, the child may note that moving the pedals brings about movement of the wheels because the pedals and wheels are connected to one another by intervening chains and gears.

Evidence that the child seeks to know about more than mere correlation or association between events comes from two studies in which children spontaneously concerned themselves with going beyond mere correlation by trying to learn, as well, about what caused one event to be correlated with another. In one study (Koslowski, in preparation) preschool children were shown an apparatus in which the movement of a bolt lock was associated with the ringing of a bell. Not satisfied with simply detecting this correlation, the majority of the children spontaneously suggested that there must be a connection someplace between the lock and the clapper and many of the children spontaneously went on to search for the connection. An even more striking

example of the concern of mediating mechanisms as a way of moving from correlation to cause comes from a study (Koslowski & Snipper, in preparation) in which turning a knob attached to a battery was associated with the ringing of an electric bell. There was a visible wire that ran from the battery to the bell. However, even when the wire was disconnected from the bell, some of the children made positive predictions that turning the knob would still cause the bell to ring. They made this prediction because they postulated an intervening connection that, though invisible, had nevertheless functioned as a mediator between cause and effect. Positive predictions were based on the premise that this intervening mechanism would continue to mediate between cause and effect even with the wire no longer contacting the bell. For example, children would suggest that the "battery stuff" or the electric stuff" from the battery would go through the wire and "shoot out" or "spray out" and hit the bell. Those children who predicted that the bell would not ring with the wire disconnected also based their predictions on considerations involving underlying mechanisms. They would argue, e.g., that the "battery stuff" would "spray out" of the wire and miss the bell. In short, even when these children could not actually see the intervening mechanism that mediated between two correlated events, they were not satisfied with merely detecting the correlation. They either tried to find the causal connection or else they postulated an invisible connection in order to explain the correlation.

The two studies just reported involved simple relationships: only one antecedent event preceded a subsequent event. In the real world, however, there are often many antecedent events that precede an effect. The child searching for an explanation in this sort of situation must first identify which particular antecedent event(s) is the causal agent before he can

concern himself with finding out what it is that mediates between or connects the antededent event (or cause) to the effect. We turn now to the issue of what sorts of cues children rely on in order to decide that one event rather than another is the cause of a phenomenon. We begin by examining how the child deals with unfamiliar situations.

When a child is faced with an unfamiliar situation (as children often are), his main task in finding a causal explanation for that situation is to somehow choose a likely cause from among a large number of possible causes. In making such a choice, we know that children are capable of relying on various rules about the way in which causes and events tend to be related. Preschool children are most likely to rely on the index of causal relationship that consists of temporal contiguity. That is, in identifying the agent that caused an effect, preschool children will search for the event that was closest in time to the effect. As children grow older, they take into account the additional cue of regularity of co-occurrence and look for an event that consistently precedes or co-occurs with the effect (Shultz & Mendelson, 1975; Siegler, 1975; Seigler & Liebert, 1974).

It must be stressed, however, that the indices of regularity and temporal contiguity are, in some sense, "cues of last resort." They are the cues that children rely on when there are no other cues present. They are cues of last resort because they do not enable one to distinguish relationships that are merely correlational from those that are genuinely causal, as well. And, we have already seen that even young children are not satisfied with simply noticing that two events are correlated: They also seek to learn about the process by which one event is able to cause the other. As we pointed out above, an important distinction between correlational and causal relationships concerns the presence of an intervening

mechanism that mediates between causally related, but not between merely correlated events. Thus, we would expect that children who are concerned with finding causal explanations would be concerned as well with information about possible intervening mechanisms that might be operating in a situation. And children, oblige us. Indeed, they oblige to such an extent that, often, even in unfamiliar situations, conclusions based on the cues of temporal contiguity and regularity will be overridden by information that suggests possible intervening mechanisms that might have mediated between cause and effect.

For example, in a study by Mendelson & Shultz (1976), two possible causal agents (i.e., antecedent events) preceded an effect. The temporal delay between one event and the effect was longer than the delay between the other event and the effect. When length of delay was the only information available, children choose as the causal agent that event which was associated with the shorter delay. That is, they relied on the cue of temporal contiguity. However, children were less likely to do this if a visible connection (in this case, a tube) mediated between the effect and that causal agent which was paired with the longer time interval. Presumably, the intervening tube suggested a rationale or an explanation for the longer delay. The children probably assumed that it took a fairly long time for the result of the causal agent to "travel through" the tube on its way to producing the effect. Thus, children were less likely to rely on temporal contiguity when they could rely instead on information about a possible intervening mechanism that could have mediated between the effect and the temporally distant event.

Just as inferences based on the index of temporal contiguity are often overridden by other considerations, the index of regularity of co-occurrence

is also less likely to be used when children can rely instead on information that suggests possible intervening mechanisms. In a study now in progress (Koslowski & Levy) preschool children are shown three instances of a particular event with each instance depicted in a separate drawing. One such event consists of a boy who has fallen from his bike. Two of the three drawings of the event also portray a particular environmental feature (i.e., a bump on the road) while all three of the drawings depict an additional environmental feature (i.e., a bumble-bee flying in the vicinity of the fallen boy). The irregularly occurring feature can be related to the event by means of a possible intervening mechanism (e.g., the bump could have caused the boy to fall by making him lose his balance). In contrast, the other feature, although it occurs regularly in all three instances, cannot be (or, at least, can not as easily be) related to the event in such a way. In spite of the fact that the presence of the bumble-bee regularly occurs with the bicycle accident, children explain the accident by citing the irregularly-occurring bump on the road as the cause. Furthermore, even when it is pointed out to them that the bump does not occur in one of the pictures, their judgments remain unshaken ("The bump was on a different part of the road that's not in the picture and the boy was able to keep his balance until he got to this part of the road.") Indeed, the main age difference in this study has to do with the number of possible intervening mechanisms that children can generate. When questioned further, many young children were unable to suggest any possible way in which the bumble-bee could have brought about the accident. Older children, in contrast, were able to suggest that, e.g., the bee (or his sting) might have distracted or scared the boy and caused him to lose his balance

that way. That is, older children were more likely to have access to collateral, or related information about the way in which the proximity of a bumble-bee could have caused the accident. This collateral information included the information that a bee might function as a distractor, that distraction can result in a loss of balance, etc. The importance of having access to collateral information will be discussed in more detail below.

We have seen that young children are capable of relying on the cues of temporal contiguity and regularity in order to decide which one of a large number of possible causes is likely to be the probable cause of an event. We have also seen that judgments based on these cues can be overridden by judgments that are based instead on information about possible intervening mechanisms (information about causal process rather than outcome). We now turn to evidence that information about the causal process can also be used in order to decide which feature (among many features of a situation) might be related to a causal explanation of an event. Thus, even when the child does not have the option of relying on the cues of temporal contiguity and regularity, he need not choose at random; he can base his choice instead on information about possible intervening mechanisms.

As an example of how this might occur in adult reasoning, consider a situation in which one is trying to explain a patient's death. One situational feature might be that the person was being treated with penicillin. On the face of it, it looks as though this feature is not causally relevant to the patient's death. Penicillin cures rather than kills. However, if we find that the patient was allergic to penicillin, then this information provides a possible intervening mechanism according to which the situational feature (treatment with penicillin) could have been causally related to the patient's death.

When confronted with simple situations, grade school children seem to engage in analogous reasoning. For example, in another study (Koslowski, et.al., in progress) children are shown a picture of a boy who has fallen from his bike. A bumble-bee is shown flying in the vicinity and a small dog is near the side of the road. When children are told that the bee was flying near the boy's face, the children are likely to incorporate the bumble-bee into a causal explanation of the event. For example, they might suggest that the bee distracted the boy or that the boy let go of the handle-bars in order to swat at the bee. Thus, they use information about the bee's proximity to the boy as suggesting an intervening mechanism by which the bee could have been related to the accident. If, on the other hand, children are told something about the bee that is causally irrelevant, (viz., that it is yellow and feels fuzzy) and are told instead the the dog was running across the road in front of the boy's bike, then children will incorporate the dog (rather than the bee) into a causal explanation of the event. In this study, as in the one reported above, an important age of difference is that older children are more able than younger ones to generate a large number of possible intervening mechanisms that could have enabled a particular causal agent (a bumble-bee or a dog) to bring about an effect. That is, they have access to more collateral information that they can rely on in order to generate a wider range of possible explanations for an event.

In the studies described in this section of the paper, we have argued that children are intrinsically motivated to detect regularity and that they also concern themselves with the process or underlying mechanism by means of which this regularity occurs. To use the terminology introduced in the first section of this paper, children seem predisposed to become initially engaged in those tasks that involve either questions about the

way the world is organized or questions about how the organization might be explained. In light of children's self-initiated concern with the process by which things happen, it is worth recalling the findings, reported earlier, that concern for process is not likely to be facilitated by a motivational system that places heavy reliance on external rewards.

Achieving Causal Explanations

Given this evidence that children show a spontaneous concern for understanding the processes or underlying mechanisms by which events are brought about, a natural question that arises is to ask about the steps or process that children themselves move through as they attempt to achieve causal explanations or descriptions of the processes by which things happen. Again, to use the terminology introduced in the first section of this paper, what can we say about the activity or manipulation that children engage in when they become involved in a task? The studies described thus far do not provide a very complete answer. In some of these studies, information about the possible intervening mechanism was actually suggested to the child by the experimenter. In other studies, although the child himself initiated a search for the intervening mechanism, the mediating connection was found so quickly that there was no time to study the search process. In order to learn more about what the search process consists of, we turn to research aimed at describing the process by which nonroutine problems are solved. The link between searching for causal explanations and solving problems is clear. Arriving at an explanation for a phenomenon is often tantamount to solving a problem. Solving a problem often requires one to learn about the underlying mechanism that connects the various parts of one another.

We will first focus on that step in the problem-solving process that consists of the exploration that precedes achievement of the correct solution. Our information is based on problems that range from the concrete to the abstract and that include: learning how to manipulate tools in order to

bring a goal-object within reach (Harter, 1930; Koslowski & Bruner, 1972; Minskaia, 1970; Richardson, 1932, 1934; Sobel, 1939; Zhukova, 1970); discovering how to compute the area of a parallelogram (Wertheimer, 1959); mastering quadratic equations (Bruner & Kenney, 1962); devising a procedure for eliminating an inoperable tumor (Duncker, 1945); and, making scientific discoveries (Wallas, 1926).

We would like to highlight five aspects of pre-solution exploration: First, pre-solution exploration must be extensive before a correct solution can be obtained (Duncker, 1945; Harter, 1930; Koslowski & Bruner, 1972; Richardson, 1932, 1934; Wallas, 1926; Zhukova, 1970). The problem-solver must be thoroughly familiar with various facets of the problem situation; superficial acquaintance will not suffice. Thorough familiarity includes learning about all aspects of the problem-situation not just those that are obviously relevant to the solution. Second, pre-solution exploration of those aspects of the situation that might constitute the means of solving the problem sometimes actually takes precedence over attainment of the goal (Koslowski & Bruner, 1972; Minskaia, 1970; Richardson, 1932, 1934; Sobel, 1939; Wallas, 1926). Pre-solution exploration is often so consuming that it becomes an end in itself -- often to such an extent that the goal or aim of the problem is actually forgotten. Notice how these first two aspects of the problem solving process would be hindered by a motivational system that overemphasized the products of learning by making rewards contingent on number of correct answers. A child in such a system would be taking a risk by exploring facets of the situation not obviously related to the correct answer or by exploring potential means to the exclusion of the goal.

The third aspect of pre-solution exploration often involves: the translation of abstract notions into concrete instances (Bruner & Kenney, 1962; Wertheimer, 1959); a reliance on analogies between the situation at hand and other situations; and, a reduction (when possible) of new problems to other,

more familiar situations (Bruner & Kenney, 1962; Wertheimer, 1959). This third aspect again brings to mind an issue that was raised earlier, viz., the importance of collateral information. Clearly, having access to a large reservoir of collateral information makes it more likely that analogies can be drawn between a current problem and other situations and also makes it more likely that there will be other familiar situations to which current problems can be seen as being similar.

An additional aspect of pre-solution exploration has to do with a possible function of this exploration. Although the evidence is not conclusive, one possibility is that exploration of the means enables the problem solver to discover which properties of the means are potentially useful in achieving a correct solution and which are largely irrelevant. For example, Sibylkin and Uzgiris (in press) had an adult model the procedure required to make an apparatus work. Some of the model's movements were, in fact, irrelevant to achieving the correct solution. During the early repetitions, children imitated both the irrelevant as well as the relevant behaviors. However, as the children came to understand for themselves how the apparatus worked, the irrelevant behaviors were no longer included in their repetitions. Zhukova (cited in Berlyne, 1970) gave preschool children (3-6 years of age) the task of bringing a lure within reach by selecting the correct one of a number of different tools -- in this case, different kinds of hooks. One group of children was given hooks of different shapes and colors; the other, hooks of different shapes but the same color. Both groups of children tried out different hooks in turn, but the second group of children achieved correct solutions faster, presumably because they did not have to explore the irrelevant cue of color as being possibly relevant to correct solution.

The final point about pre-solution exploration is that it facilitates achievement of the correct solution only if the exploring child has already achieved a certain age (the particular age being different for different problems). One can only speculate about why this is so. One possibility is that, for it to be efficacious, exploration of the means must take place in the context of previously acquired background or collateral information and that young problem solvers have not yet had a chance to acquire this background. For example, to use a rotating lever in order to solve a problem, it may well be that one needs collateral information about rotation (e.g., that direction of rotation does not matter) in order to benefit from exploration of the lever, per se.

In short, the second phase of motivated activity, the phase that involves activity or manipulation, includes a large component of exploration for its own sake, as an end in itself; a heavy reliance on collateral information; and a tendency to make use of analogies and previously mastered similar situations.

At present, we have little information about the third phase of motivated activity, disengagement. One reason for this is that, in most problem solving studies, the reason for disengagement is unambiguous: The child achieves the goal-object, whether it be a toy, the correct answer or some particular level of mastery. In many real-world situations, the goal of an activity is not as clear-cut as this. We have yet to learn what it is that makes children be satisfied with their performance of real-world tasks and their achievement of real-world goals.

Exploring the Explanations

Regarding the last phase of motivated activity, subsequent engagement, we have only slightly more information. Again, it comes from the problem solving literature and concerns the kind of exploration that sometimes

occurs after the correct solution has already been achieved. This activity consists of exploring the solution itself. For example, in a study that required children to rotate a lever in order to bring a toy within reach (Koslowski & Bruner, 1972), after children had learned how to solve the problem, they spontaneously began to investigate the solution itself by: varying the direction in which the lever was rotated; varying the way in which they positioned their hands while they held the lever; rotating the lever by pushing as well as by pulling it; etc. In a study that required children to seriate a set of size-graded sticks from smallest to largest (Koslowski, in preparation), children engaged in analogous behavior. After achieving correct seriation, some of the children spontaneously: varied the order in which sticks were chosen (largest to smallest as well as smallest to largest); varied the direction of seriation (right to left as well as left to right); and changed from keeping the bottoms of the sticks to keeping the tops arranged on a straight line. These variations were often accompanied by comments such as, "Huh! It works this way, too." Lastly, there is Piaget's observation of a child who counts, by moving clockwise, a number of elements that have been arranged in a circle. He then counts the same arrangement by moving in a counter-clockwise direction. This child also notes that it "works" both ways. We can only guess at what children are accomplishing by such exploration. One reasonable possibility is that it enables them to learn exactly which aspects of the solution are necessary, e.g., to learn that direction of rotation or direction of seriation does not alter the essential outcome. A second possibility (and one that is not incompatibe with the first) is that children use this exploration in order to fit a new-found solution into a broader context or background. For example, it may enable a child to relate his new-found knowledge of

how elements of a set can be seriated against his already acquired information about size, direction, etc. That is, it may enable him to relate new information to background or collateral information -- thus making both types of information richer and potentially more useful. Again, a system that provided extrinsic rewards for the attainment of a product would be likely not to encourage exploration of the product. Such a system would push the child to achieve, instead, yet an additional product. There could be pressure to achieve yet another correct answer rather than to more fully understand the correct answer that had already been attained.

Intrinsic Interests & Capacities in the Content of Education

We began this section of the paper with a question about the sorts of intrinsic motivation and capacities that characterize the young child. If we summarize the findings reported in this section, four general points emerge which seem to be of especial relevance to issues in education.

First, children are interested in detecting and constructing regularities or patterns in the world. In addition, they concern themselves with explanations for the underlying mechanism, or process, by which events or outcomes are brought about. In practice, this means that the initial phase of engagement is likely to be more successful to the extent that it holds the prospect of learning something about the way the world is organized or about the processes by which things happen. This also means that problems with understanding an adult's explanation or way of organizing the world may not result solely from a failure to grasp what the adult is saying. They may reflect, as well, an actual conflict between the adult's position and the one that the child has constructed. Helping the child might require first coming to understand how he has organized or explained the world on his own. Furthermore, this tendency also means that explanations

or beliefs which, from an adult's view, are wrong may nevertheless be reasonable in the sense of being warranted by the information that is actually available to (or remembered by) the child. This brings us to the second point, viz., the important role played by background or collateral information.

False beliefs about the way the world is organized or explained need not be instances of a lack of intellectual capacity or of faulty thinking. They can be, instead, what Kohler has called "good errors" -- guesses or hypotheses which, though wrong, are nevertheless reasonable. Thus, they might reflect the child's limited repertoire of (and often limited access to) the kind of factual information that would be required to achieve a correct explanation. This point is important because of its implications for education. If faulty thinking is the problem, then one type of educational remedy is called for. If, on the other hand, limited information is the culprit, then the remedy ought to consist, not of fostering new ways of thinking, but rather of providing or making available a larger body of factual information.

The third point that emerges from the above findings is the importance, in solving nonroutine problems, of exploration of the problem situation as an end in itself, without simultaneously keeping the goal or end-product in mind. We saw, in the first section, that it is exactly this sort of exploration that is undermined by extrinsic rewards. We saw, in the studies just described, not only how important such exploration is in achieving correct solutions to nonroutine problems but also that children are adequately motivated to engage in this sort of exploration and are capable of benefitting from it (assuming access to adequate background information). This suggests that periods of exploration in which the goal seems to be shunted to the side may not only be natural but may also be facilitative.