INCENTIVE AND COGNITIVE PROCESSING

Michael W. Eysenck
Birnbeck College, London

It has been found that performance on a large variety of tasks is enhanced when incentives (monetary or otherwise) are offered for high performance efficiency. Several possible theoretical explanations have been proposed to account for this effect. Edwin Locke has argued that incentives improve performance to the extent that they affect the individual's goal-setting; in contrast, Easterbrook claimed that incentives produced increased attentional selectivity with enhanced performance on the primary task being accompanied by reduced performance on a concurrent secondary or subsidiary task.

An influential theoretical position deriving from that of Easterbrook was put forward 10 years ago by Donald Broadbent. He pointed out that incentive often interacts with known arousers (e.g., white noise) in such a way as to suggest that incentive is itself an arouser. The implication is that various arousers, including incentive, are affecting the same arousal mechanism in very much the same way. In this connection, there is some evidence indicating that incentivized subjects are more physiologically aroused than control subjects.

One of the major limitations of the research in this area has been the failure to assess the effects of incentive on performance efficiency in a satisfactory manner. If incentive improves performance on a cognitive task, it does not necessarily follow that incentive has enhanced the efficiency with which all of the component processes involved in the performance of that task have been carried out. Indeed, it is even possible for an overall beneficial effect of incentive on the performance of a cognitive to mask an adverse effect of incentive on one or more of the component processes.

The cognitive task used in our ongoing research program was selected in part because it permits identification and measurement of its salient component processes. It was also selected because the effects of one arouser (white noise) on its performance have already been established by Hamilton and Hockey, thus permitting some evaluation of Broadbent's arousal-based interpretation of incentive effects. The task involves letter transformation; more specifically, the subject is given one, two, three, or four letters, and is asked to add either 2 or 4 letters to each of the original letters. Thus an example of a simple problem is 'C + 2', for which the correct answer is 'E', and an example of a relatively difficult problem is 'JEPG', for which the correct answer is 'NITK'. For each problem, the subject must work out the entire answer before responding.

What are the component processes involved in this letter-transformation task? When a letter is presented, the first process involves accessing long-term memory, locating the alphabet, and then finding the appropriate starting point within the alphabet. The second stage of processing involves the carrying out of the transformation itself, and the third stage of processing involves the storage and organization of the part-answer. In the case of a four-letter problem, these three processing stages are repeated for each letter in turn. Thus, at least conceptually, we can sub-divide the total time taken to solve a four-letter problem into 12 component stages.

A further advantage of this cognitive task is worth mentioning at this point. While it is often extremely difficult (or even impossible) to decide whether different task-processing activities occur serially or in parallel, it is virtually certain that

the component processes involved in the letter-transformation task are carried out seriatim. It is hard to see how the transformation stage could begin before the alphabet has been located in long-term memory, and it is equally difficult to believe that the answer to a letter could be stored while it is still being transformed.

In the initial study in the current series, subjects spent 5 minutes solving each of 8 different versions of the task (1, 2, 3, or 4 letters, adding 2 or 4). Incentive was a between-subjects' factor, with incentivized subjects being offered £5 (approximately 12 dollars) for obtaining an overall level of performance among the top 25% of participating subjects. Non-incentivized subjects were offered no extra monetary payment over and above their normal payment for experimental participation.

The results of this initial experiment were reasonably unequivocal. The incentivized subjects outperformed the non-incentivized subjects in each of the 8 task conditions, taking between 30% and 40% less time to solve each problem (the error rate was less than 5% in all conditions). The only significant interaction was between incentive conditions and the number of letters in each problem; this interaction involved a systematic increase in the beneficial effect of incentive as the number of letters requiring processing increased.

What do these results mean? In order to interpret the interaction between incentive conditions and number of letters per problem, we obviously need to have some understanding of the effects on the processing system of varying the number of letters. Perhaps the major effect of increasing the number of letters in the task is to increase the demands on some short-term storage system which is involved in the storage and organization of the accumulating part-answer. If so, then it may tentatively be concluded that monetary incentive increases the efficiency of a short-term storage system.

Of course, this cannot be the whole answer. Presumably one-letter problems make minimal demands on short-term storage, and yet incentive increased performance speed considerably on such problems. The implication is that incentive also affects time to access long-term memory or transformation speed (or both).

The second experiment in the series was designed to clarify the precise effects of incentive on the letter-transformation task. Only four-letter problems were used (adding 2 or 4), and the presentation was on a letter-by-letter basis. The following sequence of events occurred on each trial: the subject pressed a key in order to present the first letter; he or she then did the transformation aloud; then, the subject pressed a key in order to present the second letter; and so on. Each subject spent 40 minutes doing the task (20 minutes on each version). Incentive was manipulated as a between-subjects' factor, with incentivized subjects being offered £5 (about 12 dollars) for obtaining an overall level of performance among the top 25% of the subjects. Non-incentivized subjects only received their normal payment for attending the experiment.

The first step in the analysis of the data was to calculate 12 intra-task times. This was done by measuring the time between the first key press and the start of the transformation (assumed to reflect access time to long-term storage), the time to perform the transformation out loud (transformation speed), and the time between the end of the transformation and the next key press (assumed to reflect storage and organization). These three times were obtained for each of the four letters. The error rate was again below 5 per cent.

The major findings were quite straightforward. Incentive did not affect the time taken to access long-term storage, but did lead to increased transformation speed, especially when the add factor was 4 rather than 2. In

addition, incentive speeded up the time taken to perform storage and organization operations. This main effect was qualified by a significant interaction between incentive and letter position. In this interaction, the beneficial effects of monetary incentive were greatest during storage and organization following transformation of the third letter. In general terms, the demands on short-term storage capacity are likely to increase systematically with each additional letter. However, there is a reduced requirement for storage and organization following transformation of the fourth letter, since at that point the subject is in a position to output his or her answer to the problem. Accordingly, the interaction between incentive and letter position may be interpreted as reflecting the greater efficiency of some short-term storage system under incentive conditions.

It is interesting to compare the effects of monetary incentive and white noise on this task. Hamilton and Hockey found that noise increased the speed of transformation but decreased the speed of storage and organization. While more levels of noise and incentive must be sampled before any definite conclusions are possible, it is nevertheless interesting to note the rather different patterning of the effects of incentive and noise. In particular, incentive <u>increases</u> the speed of storage and organization, whereas noise <u>decreases</u> it. It is thus possible that noise and incentive should not be considered merely as equivalent arousing agents.

In the third experiment in the series we looked at the effects of distraction on the performance of the 4-letter, add-4 version of the letter-transformation task. The task was carried out in the presence of auditorily presented distracting stimuli (letters, numbers, or meaningless blips) which were presented on average one every 5 seconds, or in the absence of distraction. There were three within-subjects' incentive conditions: no incentive; £9 (about 20 dollars) distributed among the top 25% of subjects on low-incentive trials; and £70 (about 160 dollars) distributed among the top 25% of subjects on high-incentive trials. The session lasted approximately one hour.

One of the reasons for investigating the effects of distraction was that Easterbrook argued that incentive leads to increased concentration on task-relevant stimuli, which seems to imply that incentive should reduce distractibility. An alternative possibility is that more of the available processing resources are invested in the task under incentive conditions. If the active rejection of intermittent distracting stimuli requires processing resources, then incentivized subjects might be <u>more</u> rather than <u>less</u> distractible. A further possibility is that incentive interacts with type of distraction, so that incentive can either increase or decrease distractibility depending on the nature of the distracting stimuli. The data from the third experiment, which are currently being analyzed, will provide answers to some of these issues.

In sum, it is erroneous to assume that incentive produces an across-the-board improvement in all of the processing operations involved in the letter-transformation task. What actually happens is that simple mental operations such as those involved in transformation are speeded up by incentive, and the efficiency of some short-term store is improved. However, another processing operation (accessing long-term memory) is unaffected by incentive, perhaps because it is a relatively automatic skill. It is only by doing fine-grain analyses that one can obtain important information about the precise patterns of effects produced by incentive. The above findings have been obtained with the use of relatively modest incentives, of course. We have preliminary data suggesting that larger incentives may produce either somewhat different or very different results (as would obviously occur if one hundred thousand dollars were offered for good performance).

The Role of TAUs in Narratives

Michael G. Dyer
Computer Science Department
Yale University, New Haven CT[1]

1. Introduction

People often rely upon common sayings, or adages, when asked to characterize stories (either by way of summarization, or title selection). What are people doing in such cases? Why do adages often serve as an effective way of characterizing a story, and how are people able to accomplish this?

For instance, when asked to characterize the following story:

MINISTER'S COMPLAINT

> In a lengthy interview, Reverend X severely criticized President Carter for having "denigrated the office of president" and "legitimized pornography" by agreeing to be interviewed in Playboy magazine. The interview with Reverend X appeared in Penthouse magazine.

readers often responded with adages such as:

ADG-1: The pot calling the kettle black.
ADG-2: Throwing stones when you live in a glass house.

Clearly, these adages **are** an effective characterization of MINISTER'S COMPLAINT. But how do we recognize this fact? By what process does an 'appropriate' adage come to mind, and to what purpose?

Furthermore, when supplied with an adage and a context, some individuals experience remindings from episodes in their lives. For instance, one individual was first presented with the following:

context: EDUCATION

ADG-3: Closing the barn door after the horse has escaped.

and then asked to recall some episode from his life. He experienced this reminding:

ACADEMIA

> Years ago, I was at University U-1, where I could never get the facilities I needed for the research I wanted to do. So I decided to apply to University U-2, which offered a much better research environment. When the chairman learned I had been accepted to U-2 and was actually leaving U-1, he offered to acquire the facilities I had wanted. By then, however, my mind was already made up.

Several observations are worth making here: First, for adage ADG-3 to have initiated this reminding, the ACADEMIA episode must have somehow been indexed in long-term memory **in terms of some abstract situation** characterized by that adage. Furthermore, this indexing could **not** have had anything to do with the specific semantic **content** of the adage, since ADG-3 ostensibly concerns a farmer, a horse and a barn door. In contrast, ACADEMIA involves a chairman, a researcher, and university facilities.

To account for such phenomena, I will present a class of knowledge constructs, called **TAUs** (Thematic Affect Units), which share similarities with other representational systems under development at Yale, such as Schank's TOPs [8] and Lehnert's Plot Units [4] [5].

2. Thematic Affect Units

TAUs were first developed in the context of BORIS [3] [2], a computer program designed to read and answer questions about narratives that require the application and interaction of many different types of knowledge. In BORIS, **TAUs** serve a number of purposes: First, they allow BORIS to represent situations which are more abstract than those captured by scripts, plans, and goals as discussed in [7]. Second, **TAUs** contain processing knowledge useful in dealing with the kinds of planning and expectation failures that characters often experience in narratives. Finally, **TAUs** also serve as episodic memory structures, since they organize events which involve similar kinds of planning failures. For more detail on the use of **TAUs** in narratives, see [1].

In general, **TAUs** arise when expectation failures occur due to errors in planning. As such, they contain an abstracted planning structure, which represents situation-outcome patterns in terms of: (1) the plan used, (2) its intended effect, (3) why it failed, and (4) how to avoid (or recover) from that type of failure in the future. If we abstract out this planning structure from both the BARN-DOOR and ACADEMIA episodes, we get the following **TAU**:

TAU-POST-HOC

(1) x has preservation goal G [7] active since enablement condition C unsatisfied
(2) x knows a plan P that will keep G from failing by satisfying C.
(3) x does not execute P and G fails.
x attempts to recover from the failure of G by executing P.
P fails since P is effective for C, but not in recovering from G's failure.
(4) In the future, x must execute P when G is active and C is not satisfied.

TAU-POST-HOC captures the kind of planning failure that occurred for both the farmer who lost his horse, and the chairman who lost a graduate student. If the ACADEMIA story were told to an actual farmer who had lost his horse under the same planning circumstances, that farmer might well be reminded of his own experience. Whether this occurs or not, however, depends upon what other episodes are in long-term memory and what features are shared between them. Notice, for instance, that both BARN-DOOR and ACADEMIA share goals at some level. That is, both the farmer and the chairman had a goal requiring proximity on the part of another entity. Since these features are shared, one experience has a better chance of causing a reminding of the other to occur. For instance, the farmer would have recalled the HIRED HAND episode below before recalling the BARN-DOOR episode because of their shared features:

HIRED HAND

> The hired hand always wanted a raise, but the farmer would not grant it. Finally, the hired hand got an offer to work at a neighbor's farm. When the farmer found out, he offered the hired hand a nice raise, but it was too late.

---

[1]This work supported in part by the Advanced Research Projects Agency under contract N0014-75-C-111 and in part by the National Science Foundation under contract IST7918463.

Thanks go to Tom Wolf and Marty Korsin for helping with some of the ideas presented here, and for being sources of remindings.

Although these episodes (i.e. HIRED HAND, ACADEMIA, BARN-DOOR) share the same **TAU**, HIRED HAND and ACADEMIA have more indices in common. One possible organization for them appears below:

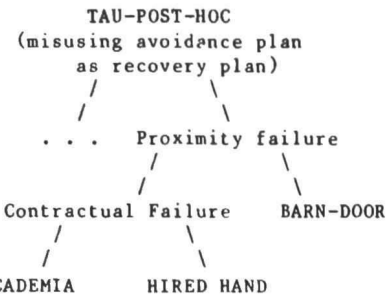

In this way **TAUs** can account for cross-contextual remindings (as in the case of BARN-DOOR and ACADEMIA). Episodes are often related in memory because they share the same abstract planning error even though they differ in content. However, cross-contextual remindings can occur only where episodes are organized under the same **TAU**, yet do not share content features. Where content is shared, the "closer" episode will be recalled.[2] Consider the following episode:

IRANIAN EMBASSY

> While holding 52 US hostages in Iran, the Iranian government condemned the take-over, by terrorists, of its embassy in Great Britain. "This is a violation of international law", protested Iran.

A reader was spontaneously reminded of this episode while reading MINISTER'S COMPLAINT (on page 1). Again, there is little in common between these stories at the content level. IRANIAN EMBASSY is about politics, while MINISTER'S COMPAINT is about pornography. However, at the abstract planning level, they both share the following **TAU**:

TAU-HYPOCRISY

x is counter-planning against y
x is trying to get a higher authority z to
  either block y's use of a plan P-1
  (or to punish y for having used P-1)
  by claiming that P-1 is an unethical plan
y claims that x has used an unethical plan P-2
  similar to P-1
therefore, x's strategy fails

In the case of MINISTER'S COMPLAINT, x is Reverend R, y is President Carter, and the third party is 'public opinion'. In the case of IRANIAN EMBASSY, x is the Iranian militants, y is the British terrorists, and the third party is 'world opinion', such as the United Nations.

As argued in [8], the reminding process is useful for this reason: Once a situation has caused one to be reminded of an episode, all of the expectations associated with that episode become available for use in making predictions about what will occur next. In the case of **TAUs**, their associated expectations include advice on either how to avoid making the error predicted by the **TAU**, or on what alternative plan can be used to recover from the error once it has been made. The ability to store **cross-contextual** episodes make **TAUs** very general and powerful mechanisms. Once an episode has been indexed under a **TAU**, its recovery/avoidance heuristics become available for use in completely different situations. Thus, planning advice learned in one context can help processing in other contexts, if the experience was recognized in terms of an appropriate **TAU** in the first place.

3. Bad Planning is Widespread

An examination of adages reveals that many are concerned with planning failures. That is, adages advice us either how to recover from a failure, or how to recognize and thus avoid future failures. Often, this advice is given implicitly, simply by describing situations in which certain planning errors lead to goal failures. In most cases, adages capture what has been called **meta-planning** [10] -- i.e. planning advice on how to select or use plans in general. For example, some adages deal with the need for checking enablement conditions before plan execution:

ADG-4 Don't count your chickens before they're hatched.

Other adages stress choosing less costly avoidance plans over more costly recovery plans:

ADG-5 A stitch in time saves nine.

or weighing the risks involved with the goal to be achieved:

ADG-6 If it ain't broke, don't fix it.
ADG-7 The cure can be worse than the disease.

Many plans require cooperation or coordination with others. This can simplify planning but complicate plan execution:

ADG-8 Two heads are better than one.
ADG-9 Two many cooks spoil the broth.

Some plans involve selecting an appropriate agent:

ADG-10 The blind leading the blind.
ADG-11 Who pays the piper calls the tune.

Timing, enablement conditions, cost, plan coordinations, and agents are just a few of the areas in which plans can go wrong. Other areas, for example, include counter-planning against a foe,

ADG-12 Cut off your nose to spite your face.

anticipating planning failures when using high risk plans,

ADG-13 Don't burn bridges behind you.

the timing of plans,

ADG-14 The early bird catches the worm.

and tradeoffs between short-term and long-term planning strategies:

ADG-15 If you can't lick 'em join 'em.
ADG-16 Don't bite the hand that feeds you.
ADG-17 Honesty is the best policy.
ADG-18 Live by the sword, die by the sword.

Any story that involves these kinds of planning failures will end up being indexed under a **TAU** which contains abstract planning advice (and can be expressed in natural language by an adage.) When a related story is read and indexed under that **TAU**, its associated adage may come to mind. For instance, a story about how a ghetto riot protesting bad economic conditions resulted in black businesses being burned, would be indexed under TAU-GREATER-HARM, with an adage such as ADG-12 possibly coming to mind.

Plans and plan failures cut across all knowledge domains. This is because we are always choosing plans, adjusting old plans to new situations, recovering from errors in planning, finding explanations for why a plan failed, etc. Furthermore, we have a large storehouse of heuristic plans, and

---

[2] This does **not** imply that BARN-DOOR can't remind one of episodes unrelated to TAU-POST-HOC. Clearly, other indexing methods may be operating at the same time. The farmer may recall BARN-DOOR simply in terms of "experiences I've had with horses". Of course, this kind of indexing can not lead to cross-contextual remindings.

there are many ways a plan can go wrong: You can't execute one plan until you have the right enablements satisfied; plan components must be executed in the right order; plans require agents, etc. This large and complex domain serves as a perfect terrain in which to index many episodes.

Many of these adages give what may appear to be superficial advice. It may seem strange that memories should be organized around such 'obvious' rules for planning, but then again, how often do we fail in our plans because we have violated some adage? How often, for instance, have we failed because we acted before we planned? ("Look before you leap.") How many times have we gotten into trouble for being late? When have we initiated a plan, only to discover we had miscalculated the amount of effort (or the side-effects) involved? ("Easier said than done.") How often have we delayed executing a simple plan, only later having to execute a more costly plan? The answer is: "very often". These adages are common because they point out the kinds of planning errors **people are always making.** By definition, plans which failed were "bad" plans. Good planners at the very least follow the general planning advice represented in the adages of their culture.

4. TAU Implementation

The recognition of **TAUs** is complex. Clearly, goals and plans must be tracked. In many cases there is also an affect component. For instance, in TAU-POST-HOC it is the futility of the recovery plan, combined with the sense of "if only I had done things differently" that helps provide an access "key" to this **TAU.**

So far the BORIS project has emphasized the use of **TAUs** in narrative comprehension. Much work remains to be done in modeling remindings during comprehension. This is important for extracting the 'moral' or point of a story. A computer program which can only answer questions of fact about IRANIAN EMBASSY, such as:

Q: How many Americans are being held in Iran?
A: Fifty-two.

Q: Who seized the Iranian embassy in Britain?
A: Terrorists.

Q: What did the Iranians do?
A: They protested the take-over.

is missing the **point** of why the IRANIAN EMBASSY is of interest. The point of IRANIAN EMBASSY is TAU-HYPOCRISY, and **that's** where it should be remembered in long-term memory (rather than just under "things I know about Iran", or "embassy events I have read").

5. TAU Experiments

What is the psychological validity of **TAUs**? Do people have **TAUs** "in their heads" and, if so, how do they use them? Some initial exploratory experiments by Seifert [9] in the Yale psychology department indicate that people use **TAUs** to organize narratives.

In one experiment, subjects read groups of stories each sharing the same **TAU**, but differing in content. Subjects were able to generate new stories, using completely different contexts, yet capturing the same planning structure specified by each **TAU.** In a follow-up experiment, different subjects were asked to sort the resulting stories generated from the first experiment. A cluster analysis [6] revealed a strong tendency for subjects to sort stories together by **TAUs.** Where stories shared the same content (but not the same **TAUs**) they were still grouped by **TAUs.**

6. Conclusions

In this paper I have presented a class of knowledge constructs, called **TAUs**, which are related to TOPs [8] and PLOT UNITs [5]. I have argued that **TAUs** organize episodes around failures in planning, and as such, **TAUs** account for at least one form of cross-contextual reminding phenomena. Furthermore, **TAUs** have adages associated with them, which express avoidance and/or recovery advice available once the **TAU** has been accessed. Since stories are indexed in terms of planning errors, this information often captures the moral or point of a story.

REFERENCES

[1] Dyer, Michael G. Thematic Affect Units and Their Use in Narratives. paper submitted to IJCAI-81, 1981.

[2] Dyer, Michael G. **In-Depth Understanding: A Computer Model of Memory for Narrative Comprehension.** PhD Thesis, Computer Science Department, Yale University, (forthcoming).

[3] Lehnert, Wendy G., Dyer, Michael G., Johnson Peter N., Yang, C. J., and Steve Harley. BORIS -- An Experiment in In-Depth Understanding of Narratives. Technical Report 188, Yale University. Dept. of Computer Science, 1980.

[4] Lehnert, W. G. Affect Units and Narrative Summarization. Technical Report 179, Yale University. Dept. of Computer Science, 1980.

[5] Lehnert, Wendy G. Plot Units and Narrative Summarization. Cognitive Science, in press.

[6] Reiser, Brian J., Lehnert, Wendy G., and Black, John B. Plot Units and the Understanding of Narratives. Cognitive Science Technical Report, Yale University (in preparation).

[7] Schank, Roger C. and Abelson, Robert. Scripts, Plans, Goals, and Understanding. Lawrence Earlbaum Associates, Hillsdale, New Jersey, 1977. The Artificial Intelligence Series.

[8] Schank, Roger C. Language and Memory. Cognitive Science 4(3), July, 1980.

[9] Seifert, Colleen. Preliminary Experiments on TAUs. unpublished manuscript. Psychology Dept. Yale University, 1981.

[10] Wilensky, Robert. Meta-Planning: Representing and Using Knowledge About Planning in Problem Solving and Natural Language Understanding. Technical Report Memo. No. UCB/ERL M80/33, Electronics Research Lab. Engineering College University of California, Berkley, 1980.