The Minimal Control Principle Predicts Strategy Shifts in the Abstract Decision Making Task

Niels A. Taatgen[1]

Invited

Abstract:

The minimal control principle (Taatgen, 2007) predicts that people strive for problem-solving strategies that require as few internal control states as possible. In an experiment with the Abstract Decision Making task (ADM task; Joslyn & Hunt, 1998) the reward structure was manipulated to make either a low-control strategy or a high-strategy more feasible. The results show that most participants are sensitive to the shifting demands of the task, consistent with the minimal control principle, but also that individual differences are large.

Keywords:

cognitive control, executive control, problem solving, decision making, individual differences

---

[1]University of Groningen

Introduction

One of the challenges of understanding human problem solving is the sheer variety of strategies that people employ to solve a particular problem. Anyone who has investigated problem solving through methods like protocol analysis can attest that it often seems as if every subject employs a different strategy (Ericsson & Simon, 1984). Problem-solving strategies from artificial intelligence are only of limited help: even though subjects show traces of behavior that resemble typical search strategies (Newell & Simon, 1972), human performance is limited by the capacities of the human cognitive architecture (Card, Moran & Newell, 1983). On the other hand, humans have a capacity that AI algorithms typically lack: the ability to adapt the strategy to new problems and changing circumstances (Taatgen et al., 2008).

An essential component of human problem solving is the ability to set certain goals and perform actions that are directed toward these goals. Goal-directed behavior requires that actions are selected on the basis of a combination of perceptual input and an internal task representation. The processes responsible for the maintenance and manipulation of an internal task representation are often referred to with the term *cognitive control* or *executive control* (Hammond & Summers, 1972). Connected to this definition, Norman and Shallice (1986) associate cognitive control with tasks that need deliberate attentional resources.

An internal task representation can encompass several types of knowledge: it entails partial results that have to maintained or carried over from one problem-solving step to the other, for example, remembering a carry in multicolumn addition, but also control information: what the current state in a complex reasoning task is. Several studies have indeed shown that these two are separable (Anderson, 2007; Borst, Taatgen, & Van Rijn, 2010). Moreover, in novice behavior deliberate control may also involve retaining and retrieving instructions for the task, further increasing the need for allocating deliberate attentional resources (Taatgen & Lee, 2003).

Cognitive control is often studied in paradigms where multiple goals have to be satisfied, for example, in concurrent multitasking (Pashler, 1994; Meyer & Kieras, 1997; Salvucci & Taatgen, 2008), in cases where the main task is interrupted (Trafton et al., 2003; Hodgetts & Jones, 2006) and has to be resumed later on, and in task switching (Monsell, 2003), where two tasks are alternated, and switching from one task to the other often incurs a cost. In most of these paradigms, cognitive control is just needed to make or keep the current task active. This is, however, not enough when tasks become more complex. A complex task requires a representation of both the task itself and the state it is in. Even though most problems afford many different solution strategies, not all are equally feasible as explanations of human behavior. Taatgen (2007) argues that people strive for a strategy that requires a minimal level of internal cognitive control. In this perspective,

control is quantified by the number of internal states that are necessary to carry out the strategy. In explicit cognitive models using cognitive architectures, it is typically easy to count the number of such states.

Minimal control has two huge advantages: it leads to the least use of cognitive resources, and it maximizes the possibility of applying the strategy to other problems. For example, in making coffee (Larkin, 1989), it is possible to employ a strategy in which each step in making coffee is planned and committed to memory. However, this method puts heavy demands on memory, and is also inflexible: if something unexpected happens during the process (e.g., the carafe is already full of water), the strategy cannot cope with it. An alternative is to have a strategy that perceives the external state of coffee making (e.g., water in carafe, filter empty), and uses this as cue to retrieve the appropriate action.

In this article, we will investigate to what extent people adhere to the minimal control strategy in the ADM task (Joslyn & Hunt, 1998), and whether they change strategies if changing task demands make a different strategy minimal. Moreover, we will investigate the hypothesis that the ability to find the minimal strategy is a source of individual differences in non-trivial tasks.

The ADM Task

ADM is a task developed by Joslyn and Hunt (1998) designed to predict individuals' performance on various real-world tasks that require decision making under time pressure. The task was used in a battery of tests designed to show individual differences in cognitive control.

ADM involves five-minute-long "games" that require subjects to classify objects into bins according to their attributes. Both objects and bins have three attributes: size, color, and shape. Only objects that match a particular bin's attributes are allowed in the bin. While objects always have three concrete attributes (e.g., small, red, circle), bins may specify any subset of the attributes (e.g., all circles). The number of attributes that a bin omits (i.e., the number of those that are wildcards) will be referred to as its generality. Thus, an "all circles" bin has a generality of two because it does not specify size or color.

Each game consists of three different bins that subjects study for as long as they want prior to starting the game. Because reviewing bins during the game is a slow and hence costly action, it is in the subject's interest to memorize the attributes of the bins. During the game, an object becomes available every 12 seconds with a pop-up notification that the subject must dismiss before continuing. None of the attributes of objects are immediately visible even when they become available. To reveal one attribute of an available object, the subject must query it by issuing the appropriate commands. To assign an object to a bin, the subject must issue a different set of commands. The subject receives points if an assignment is correct and loses points if it is incorrect. The magnitude

of reward or penalty is inversely proportional to the generality of the bin—that is, more specific bins award more points.

Because there is a time limit and objects are presented quicker than most subjects can classify them, there is a time pressure urging subjects to act as quickly as possible.

For example, the bins in the first game are:

Bin 1: Tiny red circles (150 points)

Bin 2: Red objects (50 points)

Bin 3: Medium squares (100 points)

Objects in the game are initially unknown, so subjects have to query the system about the attributes of each object. For example, the program might announce that object #1 is available, after which the subject has to decide what to do with it. For example, she might decide to ask for the color of object #1, to which the answer could be "red." Based on that knowledge, the object could be sorted into Bin 2 right away (yielding 50 points), or the subject might decide to inquire about other attributes to be able to sort the object in a more specific and therefore more profitable bin. Even without time pressure, the strategy to find the bin with the maximum score is far from trivial. For example, if it is known that an object has size medium, it still not clear whether it would fit in Bin 2, 3, or both. On the other hand, if the shape of an object is triangle, it can only go into Bin 2. The time pressure in the experiment presents additional challenges: if we know the object is red, is it worthwhile to inquire both other attributes in order to check whether it fits into the profitable Bin 1, or is it better to just assign it to Bin 2? Or shall we just assign it to Bin 1 and hope for the best?

Dickison and Taatgen (2007) determined on the basis of an earlier ADM experiment that subjects employ two types of strategies, which they characterized as high control and low control. The high-control strategy entails retrieving bins from memory, one at a time, and checking them against the known attributes. If this yields a unique winner, the object is assigned to that bin, otherwise more attributes are queried. The low-control strategy involves only a single memory retrieval of a bin based on the currently known attribute values. This retrieval relies on associative links in memory to retrieve an appropriate candidate. If carried out correctly, the high-control strategy always leads to correct assignments, while the low-control strategy does not. The high-control strategy requires a fairly complex mental representation, because there are several possible situations that have to be handled differently (no candidate bin has been identified, a candidate has been retrieved that has been rejected, a candidate has been retrieved that is potentially overly specific so other candidates have to be checked, a candidate has been retrieved that is overly general, so more specific candidates have to be checked, etc.) The low-control strategy requires far less executive control, because it just entails retrieving a single

bin from memory. It is therefore generally faster, but may lead to wrong or suboptimal assignments.

In the example discussed above in which it has already been determined that the size of the object is medium, the low-control strategy would retrieve a bin with "medium" as a cue, retrieving Bin 3, because that bin is the only one with that attribute. It would then assign the object to the bin, even though this may be incorrect (e.g., when it is a medium red triangle). The high-control strategy would also retrieve Bin 3 first, but would then conclude that the object could not be assigned yet without further checks. It would then retrieve Bin 2, and conclude that the object also fits that bin, and therefore query more attributes. In the case of a triangle object the low-control strategy would query more attributes, because a retrieval based on triangle would not produce anything, while the high-control strategy would figure out that Bin 2 is the only candidate.

Although the choice of strategy may be a source of individual differences in the sense that the cognitive demands of the high-control strategy may be too high for some, the demands and rewards of the task environment may also impact the choice. Here a slightly different source of individual differences may show up: the ability to optimally adapt the task strategy to the current environment. To investigate this, we conducted an experiment in which the environment encouraged subjects to use a low-control strategy, at least initially. In the first few games, there is no penalty for assigning an object to the wrong bin. The low-control method is therefore reasonably feasible, because if the object is assigned to the wrong bin, it can be assigned to another bin in the next step. Two other factors favor a low-control strategy: every 12 seconds a new object is introduced, introducing time pressure into the game, and the arrival of a new object is announced by an interruption that the subject has to dismiss in order to carry on. Moreover, after the interruption the program continues with the main menu, which means that any previous selections that the subject has made are gone.

After the initial games, though, the environment changes: now penalty points are given for wrong assignments (subjects are told this at the beginning of the game), and the experiment rigs the odds of guessing to the extent that if the subject tries to guess a bin while there are multiple candidates, he almost always guesses wrong (subjects do not know this). In the last two games the environment is switched back to the initial one.

Our expectation for the experiment is that most subjects will opt for a low-control strategy in the first games of the experiment, but that they switch to a high-control strategy when the environment changes. Once the environment changes back, we expect subjects to switch back to the low-control strategy. There are several ways to assess the kind of strategy subject employ. When a subject assigns an object to a bin, we can check whether this assignment is justified by the queries the subject has made. The high-control strategy should correspond with a justification score close to 1, because the subject only assigns objects to bins once that assignment is justified. The low-control strategy

should yield a much lower justification score. Another variable to look at is the number of queries subjects make: in most cases the high-control strategy requires more queries. A third variable to look at is whether the bin the subject chooses is optimal, that is, the bin with the lowest generality that still fits the object (e.g., putting a tiny red circle in Bin 2 is justified but not optimal). Although interesting, the optimality variable does not necessarily distinguish between the two strategies. A low-control strategy may involve trying the most specific bin first, yielding a high optimality score, or may leave that to chance, with a more average optimality score. Similarly, a high-control strategy might query the system until the optimal solution is found, or it might decide to assign an object to a bin as soon as it is justified, potentially sacrificing optimality.

Method

Participants

Thirty-three subjects were recruited from the Carnegie Mellon University campus and its surrounding community. All experiments were run with the consent of participants and with approval of the Carnegie Mellon IRB. Subjects were paid $10 for their involvement.

Materials

The task consisted of 1 practice game followed by 6 real games, for a total of 7 games. All games lasted 5 minutes. All games involved 3 bins with a presentation rate of 1 object every 12 seconds, except for the practice game in which the presentation rate was 1 object every 15 seconds. Bins in each game were different, but always with one bin of each of generality 0, 1, and 2, yielding 150, 100, and 50 points, respectively.

Procedure

Participants were instructed to read through four pages of instructions presented on paper. These instructions contained several screen shots to familiarize the subjects with the interface. The instructions detailed the objective of the game and how to query objects and assign objects to bins, as well as the point structure whereby more specific bins yield higher points than more generic ones. An example procedure for querying features of an object and assigning it to a bin is also provided in the instructions. The instructions also mentioned that in the first set of games there would be no penalty for making wrong assignments, and pointed out that trying out several bins is a possible strategy (to make sure that the subjects know that this is not considered cheating). Participants were free to ask questions after reading the instructions.

For each game, a preview phase preceded the timed portion of the game. During the preview phase, participants could review the bins available for that game for as long

Figure 1. Screenshots of the experiment. Left panel shows the main menu, and the right panel shows a menu in which an object has to be chosen to be queried.
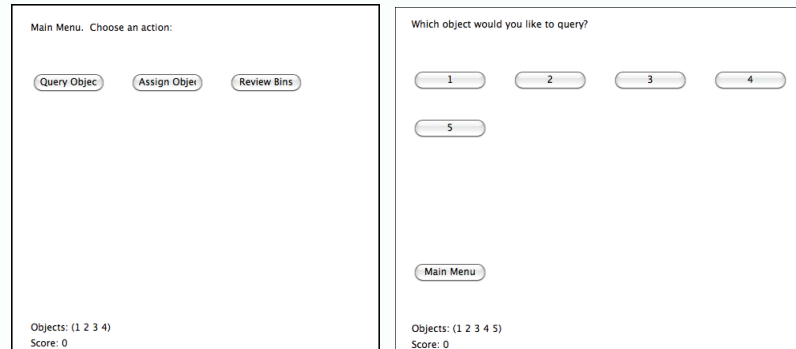
as desired. The features that each bin accepts as well as the number of points each bin yields were presented. When ready to begin, participants would press the "start" button to begin the timed portion.

The game progressed via a menu-based system in which menu choices were made using buttons (see Figure 1 for two examples). At the top-level menu, participants could choose to query an object feature, assign an object to a bin, or review a bin's attributes list. The subjects were required to make at least one query for every object. To query an object feature, after clicking "query" button, participants had to select an object from an object menu (which would contain a button for every object number that wasn't assigned to a bin yet), and then one of the words "color," "shape," or "size" from the attribute menu. After the subject had selected one of the attributes, the interface would display the value of that attribute, for example "Object 1 is RED." To assign an object to a bin, participants were prompted to enter an object number followed by a bin number. Feedback was then presented showing whether the assignment was successful, and how many points were awarded or penalized. A running total of the current score was shown at all times along the bottom of the screen, as well as a list of currently available object numbers.

In addition to the list of available object numbers, as new objects became available, a "pop up" prompt was shown, interrupting the participant's interaction with the computer. When a pop up showed up, participants were required to click a button to acknowledge the notification, which would bring them back to the main menu.

In the practice game and in Games 1, 2, 5, and 6 there was no penalty for making a wrong assignment, and it was possible to try a different assignment with the same object. When the subject made an assignment while the queried attributes of the object still allowed multiple assignments, the program would randomly pick one of the remaining candidate bins and determine the yet unqueried attributes to fit that bin. In other

words, the identity of the object was not predetermined, but depended on the choices of the subject. In Games 3 and 4 penalty points were given if the object was assigned to a wrong bin. These penalty points were –150, –100, and –50, for the bin of generality 0, 1, and 2, respectively. Moreover, if the subject assigned an object to a bin without justification, there was a 90% probability that the program would set the objects' attributes such that it would fit one of the other candidate bins (so, assigning a medium object to Bin 3 would have a 90% probability of only fitting Bin 2 and not Bin 3, basically implementing Murphy's Law). All these factors strongly favor a high-control strategy.

Results

Global Results

In most analyses we will compare stages 1, 2, and 3 of the experiment. Stage 1 includes Games 2 and 3, stage 2 Games 4 and 5, and stage 3 Games 6 and 7. In stage 1 and 3 the environment favors a low-control strategy, while stage 2 favors a high-control strategy.

In each game, subjects can assign a maximum of 25 objects. Figure 2 shows that subjects are on average close to this maximum. In stage 2 the number of assigned objects does decrease due to the increased difficulty, but not by many. This decrease is significant though, as an analysis of variance reveals: $F(2,64) = 20.7$, $p < 0.001$, $\eta^2 = 0.39$. There is a difference both between stage 1 and 2 ($t(32) = 4.6$, $p < 0.001$), and between stage 2 and 3 ($t(32) = 4.9$, $p < 0.001$).

For the purpose of the analysis we consider each assignment to a wrong bin as an error. Figure 3 shows the mean error rates for each of the games in the experiment.

The graph clearly suggests that the error rate decreases in the high-penalty games,

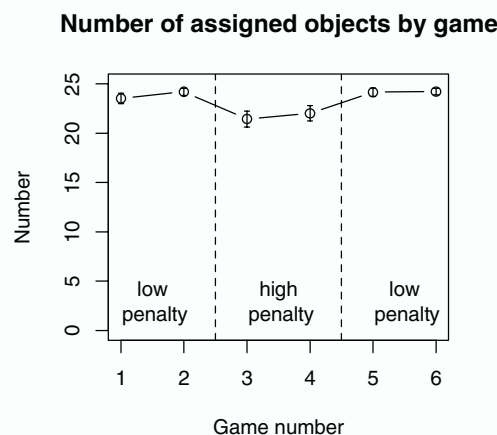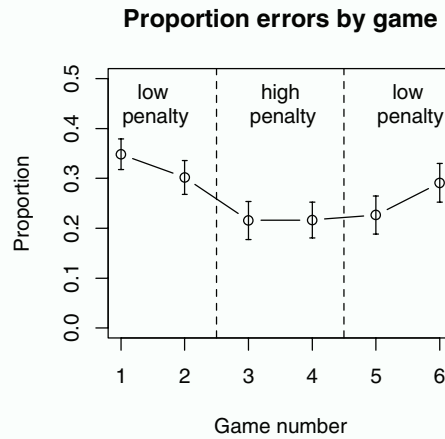**Figure 2.** Number of objects assigned by game. Error bars in this and subsequent figures show 1 standard error.

Figure 3. Error rates for each of the six games.

Proportion errors by game

and gradually increases again with the later low-penalty games. To analyze these results we averaged the error rates for the three stages of the experiment, and entered these into an analysis of variance with stage as independent variable. The results show an overall effect of stage, $F(2,64) = 5.80$, $p = 0.005$, $\eta^2 = 0.15$, with stage 1 different from stage 2, $t(32) = 3.31$, $p = 0.002$, but stage 2 not different from stage 3, $t(32) = 1.5$, $p = 0.14$.

A more precise assessment of strategy can be determined on the basis of whether a particular assignment of an object to a bin is justified. This means that on the basis of the available information the object can be assigned to the bin, even though there might be a more profitable bin to which it could also be assigned.

Figure 4 shows the proportion in which the first assignment to a bin was justified. Clearly, this rate increases in the high-penalty games and decreases again in the later

Figure 4. Proportion of the assignments that were justified.

Justified by game
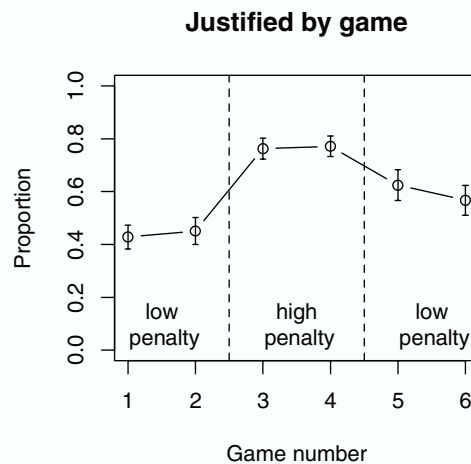
low-penalty games. The analysis of variance confirms this; there is an effect of stage on justification rate, $F(2,64) = 23.2$, $p < 0.001$, $\eta^2 = 0.42$, and both a change between stage 1 and stage 2, $t(32) = 7.5$, $p < 0.001$ and stage 2 and stage 3, $t(32) = 3.48$, $p = 0.001$.

The high-control strategy generally requires more queries to determine a choice, so we expect the number of queries to go up in the high-penalty games.

Figure 5 shows that the number of queries indeed increases sharply. An analysis of variance confirms the impact of stage on the number of queries, $F(2,64) = 15.5$, $p < 0.001$, $\eta^2 = 0.33$. Subsequent tests reveal that the number of queries indeed goes up between stage 1 and stage 2, $t(32) = 6.5$, $p < 0.001$, and down again between stage 2 and stage 3, $t(32) = 2.70$, $p = 0.011$.

The final measurement that we consider is the proportion of optimal choices, that is, the rate at which the bin with the highest point value is selected.

Figure 6 shows that there is not much change in the proportion of optimal choices, which is confirmed by an analysis of variance, $F(2,64) = 2.08$, $p = 0.13$, $\eta^2 = 0.06$. This lack of change is remarkable, given the changed circumstances of the task. Nevertheless, this null-result was to be expected on the basis of the two strategies. The high-control strategy aims at finding the correct and optimal choice, so a high score on optimal is no surprise. The low-control strategy tends to retrieve the more specific bins, and therefore typically picks optimal (but not necessarily correct) choices as well.

The overall picture that emerges from the averaged results is that subjects indeed opt for an initial low-control strategy, and then switch to a high-control strategy in stage 2. When the experiment switches back to a configuration that favors the low-control strategy, participants seem to revert to the low-control strategy, but not back to the levels of stage 1.

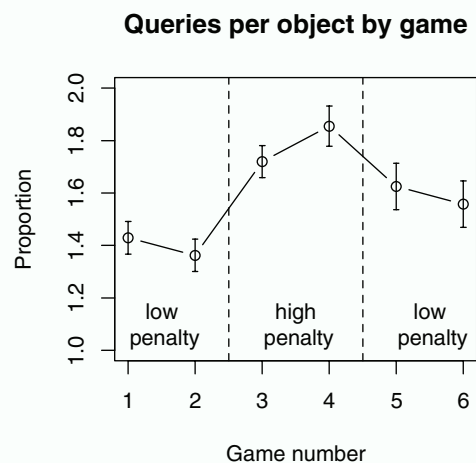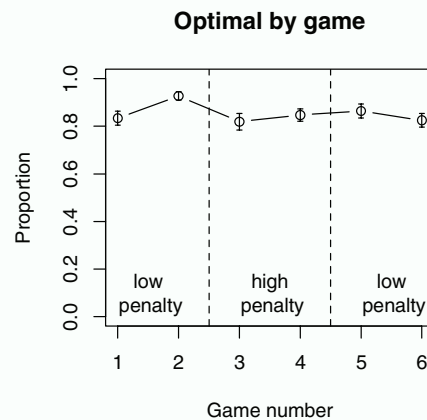**Figure 5.** Mean number of queries per object by game.

Queries per object by game

Figure 6. Proportion optimal choices per game.

Optimal by game

Individual Differences

Even though the mean results show a general pattern of adapting strategy to environment, this does not mean that every subject follows that pattern. Indeed, inspection of individual differences shows that subjects can follow several trajectories. Figure 7 shows eight representative examples of the proportion of justified assignments for the six games. Subjects 27 and 30 show the clearest adaptation to the changing circumstances. In fact, Subject 27 had the highest overall score of all subjects. Other subjects are slower to adapt: Subject 13's performance suggests a justification rate that lags behind the environmental changes. Subject 3 switches from a low- to a high-control strategy, but does not change back in Games 5 and 6. Subject 5 only exhibits a very minor adaptation to changing circumstances. Subjects 4, 9, and 14 show no adaptation at all, even though they are quite different. Subject 4 has the lowest score of all subjects, and Subject 14 the second lowest. Subject 9 only ranked 19 (out of 33), even though he or she always made assignments that were fully justified.

Is switching strategies based on the circumstances optimal in order to obtain a maximum score? The fact that the highest-scoring subjects do so is already an indication. As further evidence we looked at the correlation between the proportion of justified bin assignments and overall score.

Figure 8 shows that in Games 3 and 4 the proportion justified assignments strongly correlates with score. However, in the other games there is only a small (non-significant) negative correlation, showing that the low-control strategy is at least as good if not better than the high-control strategy in those games.

A final question we have to ask is whether low-control and high-control are indeed separable strategies. The strong shifts in proportion justified bin assignments in individuals

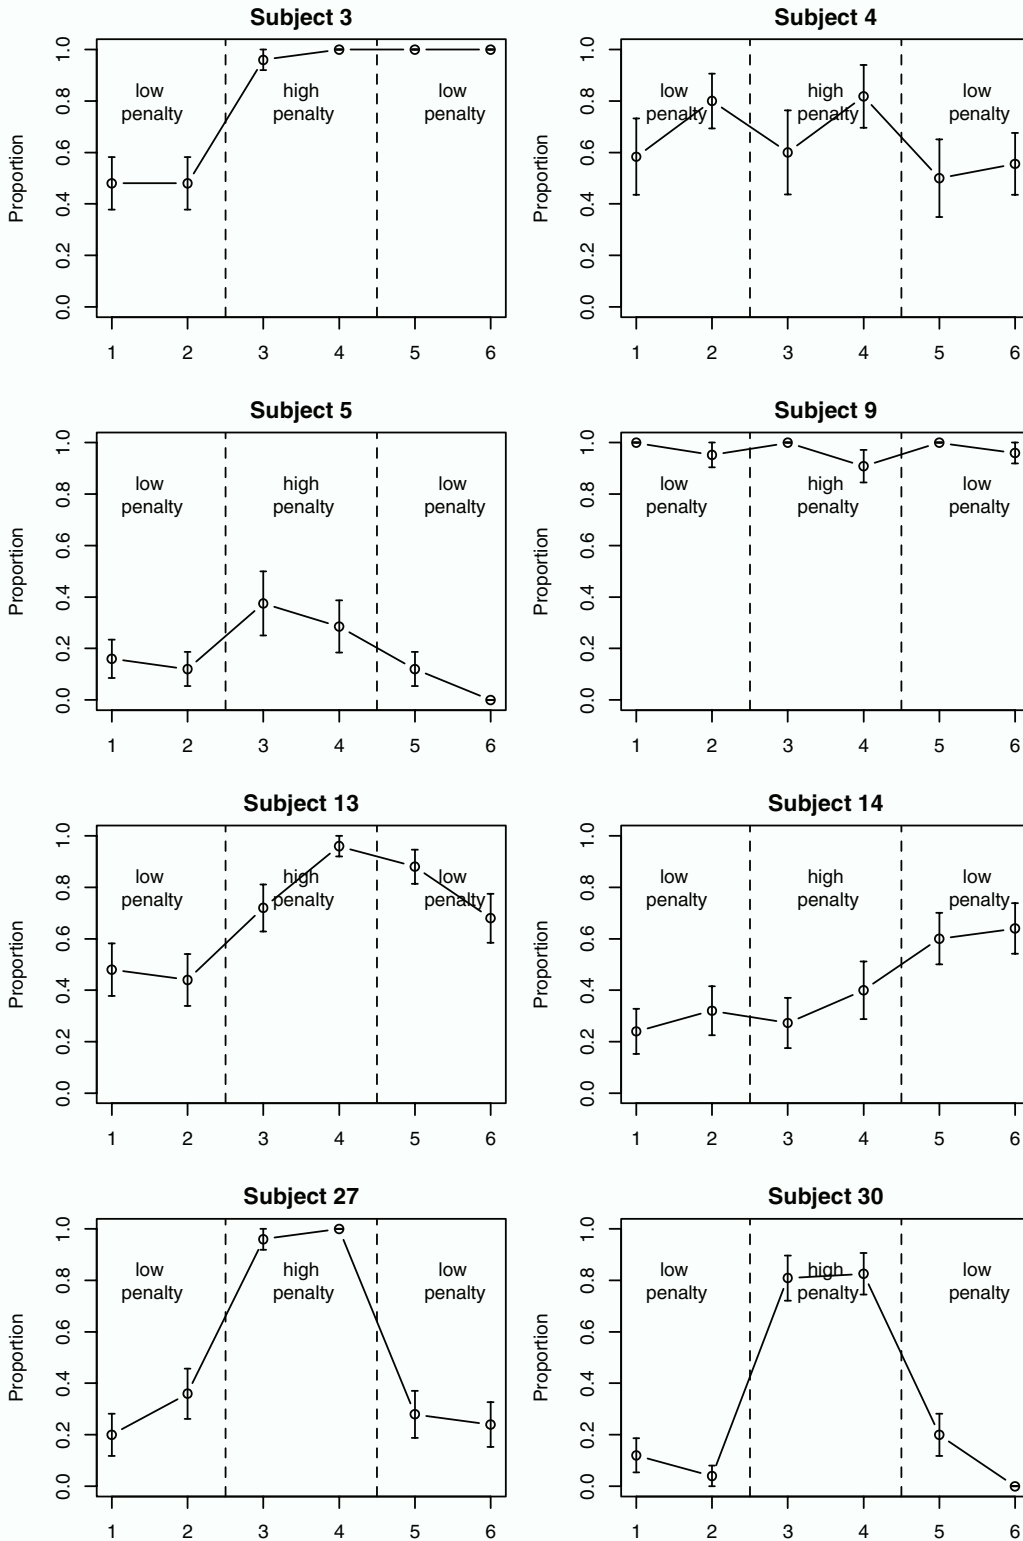**Figure 7.** Proportion of justified bin assignment for eight subjects by game.

Figure 8. Scatterplots of the relationship between the proportion of justified bin assignments and overall score.
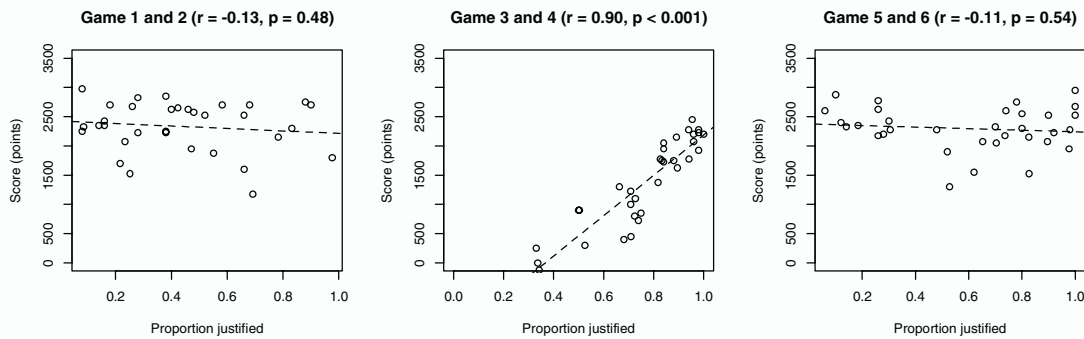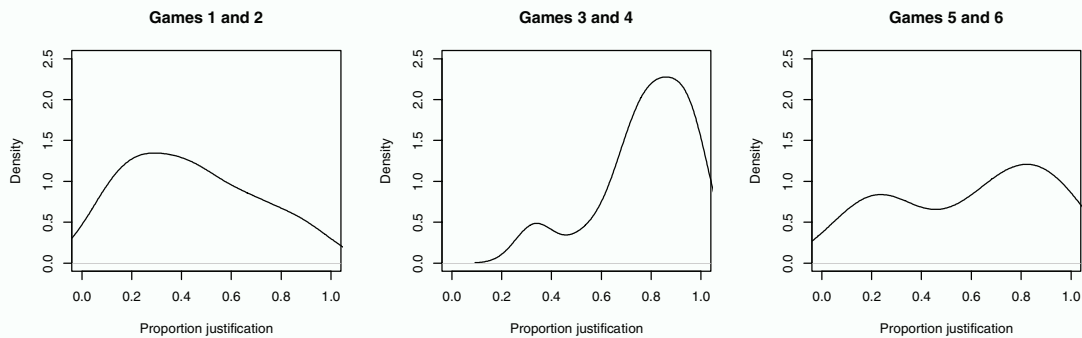

Figure 9. Density plots of the proportion justified bin assignments in the three stages of the experiment.

(see Figure 7) suggest there is. Nevertheless, a true bimodal distribution is hard to demonstrate. Figure 9 shows density plots of the three stages of the experiment. The figure shows that there is at least some indication that the distribution has two peaks.

General Discussion

The experiment has demonstrated that people adjust the complexity of their problem solving to match the task demands. In the first stage of the experiment a low-control strategy was sufficient to perform well on the task, which prompted most subjects to adopt such a strategy. In the second stage, both the implicit and explicit reward structure demanded a more elaborate high-control strategy, which most subjects adapted. In the third stage the low-control strategy was again sufficient for good performance. Although many subjects shifted back to a low-control strategy some persisted with the high-control strategy. The

results generally support the minimal control principle (Taatgen, 2007), which states that the smallest amount of control (measured in control states in a model) that still performs well on the task is the most likely accurate representation of human performance.

Even though most subjects were responsive to the changes in task demands, some of them were not. Indeed, the ADM task is designed to explore individual differences in executive control (Joslyn & Hunt, 1998). For some, the high-control strategy may exceed certain cognitive capacities, in particular, working memory capacity to process multiple bins instead of just one. But for others, the balancing act of finding minimal control while still maintaining performance may be too challenging. Subjects who do not return to the low-control strategy in stage 3 may be an example of that. On the other hand, the high-control strategy does not lead to lower performance at this stage (or at least not much), it just requires more mental effort.

Looking at individual differences from a perspective of strategic adaptation is relatively unusual in the domain of computational models of cognition. A more typical approach to model individual differences is to propose some parameter—such as working memory capacity—that varies among individuals and observe how changes in that parameter affect the model's performance (Daily, Lovett, & Reder, 2001; Taatgen, 2002; Rehling et al., 2003; Chuderski, Stettner, & Orzechowski, 2006). Varying parameters is, however, not the solution to all individual difference problems, and would not have been appropriate for this experiment.

An essential aspect of the ADM task is that the steady arrival of new objects produces interruptions. A low-control strategy makes it easier to recover from interruptions, because it requires a smaller mental representation of the problem (Borst, Taatgen, & van Rijn, 2010; Trafton et al., 2003). Moreover, the low-control strategy is faster, making it possible to assign an object to a bin before the next one arrives.

In summary, the ADM task and experiment show that human problem solving is constrained by cognitive demands but adapts well to changes in demand of the environment. The ability to adapt well to changing circumstances appears to be a source of individual differences that can be indicative for cognitive control.

References

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press.

Borst, J. P., Taatgen, N. A., & Van Rijn, H. (2010). The problem state: A cognitive bottleneck in multitasking. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 36*(2), 363–382.

Card, S. K., Moran, T. P., & Newell, A. (1983). *The Psychology of Human-Computer Interaction*. Erlbaum.

Chuderski, A., Stettner, Z., & Orzechowski, J. (2006). Modeling individual differences in working memory search task. *Proceedings of the Seventh International Conference on Cognitive Modeling* (pp. 74–79).

Daily, L. Z., Lovett, M. C., & Reder, L. M. (2001) Modeling individual differences in working memory performance: A source activation account. *Cognitive Science, 25*, 315–353.

Dickison, D., & Taatgen, N. A. (2007). ACT-R models of cognitive control in the abstract decision making task. *Proceedings of the Eighth International Conference on Cognitive Modeling* (pp. 79–84). Psychology Press.

Ericsson, K. A., & Simon, H. A. (1984). *Protocol Analysis. Verbal Reports as Data.* MIT Press.

Hammond, K. R., & Summers, D. A. (1972). Cognitive control. *Psychological Review, 79*(1), 58–67.

Hodgetts, H. M., & Jones, D. M. (2006). Interruption of the tower of London task: Support for a goal-activation approach. *Journal of Experimental Psychology: General, 135*(1), 103–115.

Joslyn, S., & Hunt, E. (1998). Evaluating individual differences in response to time-pressure situations. *Journal of Experimental Psychology: Applied, 4*(1), 16–43.

Larkin, J. H. (1989). Display-based problem solving. In D. Klahr & K. Kotovsky (Eds.), *Complex Information Processing: The Impact of Herbert A. Simon* (pp. 319–341). Erlbaum.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: I. Basic mechanisms. *Psychological Review, 104*(1), 3–65.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences, 7*(3), 134–140.

Newell, A., & Simon, H. A. (1972). *Human Problem Solving.* Prentice-Hall.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and Self-regulation* (pp. 1–18). Plenum Press.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin, 116,* 220–244.

Rehling, J., Demiral, B., Lebiere, C., Lovett, M., & Reder, L. M. (2003). Modeling individual difference factors in a complex task environment. In F. Detje, D. Doerner, & H. Schaub (Eds.), In *Proceedings of the Fifth International Conference on Cognitive Modeling* (pp. 287–288). Universitats-Verlag Bamberg.

Salvucci, D. D., & Taatgen, N. A. (2008). Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review, 115*(1), 101–130.

Taatgen, N. A. (2002). A model of individual differences in skill acquisition in the Kanfer-Ackerman Air Traffic Control Task. *Cognitive Systems Research, 3*(1), 103–112.

Taatgen, N. A. (2005). Modeling parallelization and flexibility improvements in skill acquisition: from dual tasks to complex dynamic skills. *Cognitive Science, 29*(3), 421–455.

Taatgen, N. A. (2007). The minimal control principle. In W. Gray (Ed.), *Integrated Models of Cognitive Systems*. Oxford University Press.

Taatgen, N. A., Huss, D., Dickison, D., & Anderson, J. R. (2008). The acquisition of robust and flexible cognitive skills. *Journal of Experimental Psychology: General, 137*(3), 548–565.

Taatgen, N. A., & Lee, F. J. (2003). Production compilation: A simple mechanism to model complex skill acquisition. *Human Factors, 45*(1), 61–76.

Trafton, J. G., Altmann, E. M., Brock, D. P., & Mintz, F. E. (2003). Preparing to resume an interrupted task: effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies, 58*, 583–603.

Acknowledgments