

Distraction and action slips in an everyday task: Evidence for a dynamic representation of task context

MATTHEW M. BOTVINICK and LAUREN M. BYLSMA
University of Pennsylvania, Philadelphia, Pennsylvania

We report here a novel and counterintuitive effect of distraction on routine sequential action. The effect, predicted by a recent computational model of sequential behavior, relates to the tendency for a momentary distraction, such as a brief interruption, to lead to subsequent slips of action. The specific prediction is that errors should be more likely following a distraction occurring toward the middle of a subtask sequence than following a distraction occurring at the end of a subtask. This was tested and confirmed in an experiment involving repeated performance of an everyday task (coffee making) under conditions involving frequent interruption. The observed effect provides differential support for existing models of sequential behavior and offers a highly constraining benchmark for future theories.

Making the bed, taking a shower, getting dressed . . . From the moment we wake up each day, a large portion of our time is occupied with familiar, goal-oriented routines. For the most part, such tasks are completed with ease and according to plan, a testimony to the robustness and flexibility of human information processing. Nevertheless, even in the most familiar of tasks, errors or slips of action do occur at nontrivial rates. Since William James (1890), such errors have attracted the interest of psychologists (see, e.g., Baars, 1992; Reason, 1990), and more recently, they have become a focus in human factors research (e.g., Duffy & Saull, 2003). At least two considerations make it seem important to understand their nature and their origins. First, depending on the setting, even minor slips can have serious repercussions, as, for example, in the case of aircraft operation. Second, the characteristics of action slips may provide insight into the mechanisms that underlie routine sequential behavior at large (Heckhausen & Beckmann, 1990; Schwartz, 1995).

In the present article, we describe a previously unreported behavioral effect pertaining to everyday slips of action. The effect, which concerns the impact of momentary distraction on the frequency of action slips, was predicted by a recently proposed theory of naturalistic action and places a new constraint on theoretical models of both action slips and sequential behavior in general.

Slips, Decision Points, and Attentional Checks

Influential theories concerning the origins of action slips have been proposed by Norman (1981) and Reason

(1990; Reason & Mycielska, 1982). One important observation shared by both of these accounts is that slips of action often occur at *decision points*, junctures where selection of the appropriate action cannot be accomplished solely on the basis of the immediately preceding action or the state of the environment but, instead, relies on accessing information about the broader task context, including earlier actions and previously established goals. An example of such a decision point can be drawn from the familiar task of preparing coffee with cream and sugar. In this task, as in many others, decision points occur at the transitions between subtasks¹—in particular, the transition points coming at the end of cream adding and sugar adding. Consider, for illustration, the situation at the end of cream adding. The correct next action here depends on whether sugar has or has not already been added. If it has not, the appropriate next step is to locate the sugar and to begin the procedure for adding it to the coffee. If sugar has been added, the task is complete, and another activity, such as drinking the coffee, may begin. Selecting the correct action at this decision point thus requires accessing information about actions already completed or goals already accomplished.

Why do slips tend to occur at such decision points? According to Reason (1992), the demands of action selection at these points require the actor to enter a special attentional mode. Norman (1981) makes a similar proposal, referring to these junctures as “attentional checkpoints.” According to both theories, decision point errors arise due to a failure to perform the special attentional operation, or check, that ordinarily ensures task-appropriate behavior. Thus, for example, the error of adding sugar to one’s coffee, after having already added sugar and then cream, would be attributed to a failure to execute an attentional or mnemonic operation, which would ordinarily have established whether sugar had yet been added.

The present work was supported by National Institutes of Health Award MH 16804 (M.M.B., PI). Correspondence concerning this article should be addressed to M. M. Botvinick, Center for Cognitive Neuroscience, University of Pennsylvania, 3720 Walnut Street, Philadelphia, PA 19104-6241 (e-mail: mmb@mail.med.upenn.edu).

An Alternative Account

In recent work, Botvinick and Plaut (2003, 2004) have presented a model of routine sequential action that provides a rather different account of action slips and, in particular, their tendency to occur at decision points. The model, which is illustrated in Figure 1, takes the form of a recurrent neural network, mapping from perceptual inputs to action outputs via a set of internal units. As a result of its recurrent connectivity, the network maintains, over its internal units, a distributed representation of temporal and task context. Applying this model to the everyday routine of coffee making, Botvinick and Plaut (2003, 2004) demonstrated its ability to account for a variety of benchmark phenomena relating to normal error-free behavior, apraxic performance, and (most relevant here) slips of action.

As in the empirical data, errors in the Botvinick and Plaut (2003, 2004) model tend to fall at decision points—that is, the transitions between subtasks. However, such errors are not due to a failure to execute a checking operation. Instead, decision point errors occur because the system's online representation of task context has become degraded. Thus, for example, the error of adding sugar to coffee twice occurs not because a check has been omitted, but because disruptions to the model's internal representations have resulted in a loss of the information that sugar had been added earlier.

Importantly, according to this model, the degradation that leads to decision point errors can occur at any time, not only at the boundaries between subtasks. Indeed, a distinctive claim of the Botvinick and Plaut (2003, 2004) model is that context information is most susceptible to loss toward the *middle* of subtask sequences. The explanation for this relates to how differences in temporal or task context are internally represented. In the model, distinctions between differ-

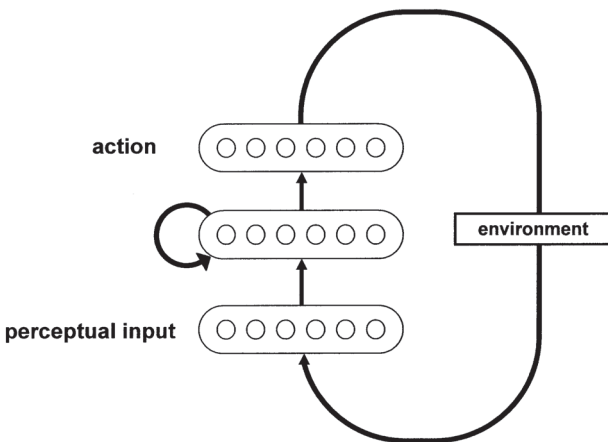


Figure 1. Schematic representation of the neural network model proposed by Botvinick and Plaut (2003, 2004). Arrows indicate full connectivity—that is, every unit in the sending layer has a connection to every unit in the receiving layer. The implemented model contained 39 input, 50 internal, and 20 output units. From “Doing Without Schema Hierarchies: A Recurrent Connectionist Approach to Normal and Impaired Routine Sequential Action,” by M. Botvinick and D. C. Plaut, 2004, *Psychological Review*, 111, p. 400. Copyright 2004 by the American Psychological Association. Adapted with permission.

ent contexts are represented very robustly close to decision points, where such distinctions are directly relevant to action selection. However, elsewhere—and in particular, toward the middle of subtask sequences—differences in temporal context are represented less strongly. This, in turn, makes it easier for the system's representation of context to become disrupted, setting the scene for a later decision point error.

The point can be illustrated by again considering the decision point falling at the end of adding cream to coffee. As has already been noted, the critical piece of context information needed here is whether or not sugar has been added. In order for this information to be available, it must be preserved throughout the course of the cream-adding subtask. However, note that the information is not actually *used* until the decision point falling at the end of the subtask. It is not directly relevant to action selection at the steps belonging to cream adding itself, since cream adding is executed in the same fashion regardless of whether sugar has been added. According to the Botvinick and Plaut (2003, 2004) model, this lack of immediate relevance leads the context information to be represented relatively weakly toward the middle of the subtask (see Figure 2). This, in turn, renders it more susceptible to disruption and loss during performance of mid-subtask actions.

This aspect of the Botvinick and Plaut (2003, 2004) theory leads it to make a distinctive prediction about the effect of momentary distraction. Earlier, it was suggested that according to the attentional check theory, distraction should be most disruptive when it occurs at end-subtask, just before a decision point, since it is here that distraction is most likely to interfere with execution of the checking operation. In contrast, the Botvinick and Plaut model predicts that distraction should be most disruptive when it falls toward the middle of a subtask sequence, even though the errors that result from such distraction do not occur until the end of the subtask.

The prediction can be illustrated on the basis of a simulation reported by Botvinick and Plaut (2004). Here, the impact of momentary distraction in the coffee-making task was simulated by injecting noise into the model's unit activations on a single cycle of processing. Figure 3 displays the relevant simulation results. When distraction occurred at the end of the cream-adding subtask, immediately preceding a decision point, few if any errors resulted. However, when the same intervention was made at an earlier, mid-subtask step, significantly more errors occurred at the subsequent decision point.

The Present Experiment

In order to test the predictions of the Botvinick and Plaut (2003, 2004) theory, we studied the performance of normal participants on an everyday task under conditions involving intermittent distraction. In keeping with the work that generated the prediction, we employed the task of coffee making. Distraction was imposed by momentarily interrupting the subjects' performance on this task, requiring them to perform a secondary arithmetic task. Interruptions were timed to occur either at the end of a subtask (i.e., just before a decision point) or at mid-subtask. The impact of




Figure 2. Attenuation of context information toward the middle of a subtask sequence in the Botvinick and Plaut (2004) model. Numbers on the x-axis indicate steps in the cream-adding subtask (1. locate creamer, 2. put down spoon, 3. pick up creamer, etc.). At each step, the specific pattern of activation over the model's internal units depends on whether sugar has or has not already been added. The point plotted at each step indicates the distance, within the model's 50-dimensional representational space, between the patterns arising in these two contexts. The U-shaped pattern reflects the attenuation of context information toward the middle of the subtask. From "Doing Without Schema Hierarchies: A Recurrent Connectionist Approach to Normal and Impaired Routine Sequential Action," by M. Botvinick and D. C. Plaut, 2004, *Psychological Review*, 111, p. 412. Copyright 2004 by the American Psychological Association. Adapted with permission.

these two forms of distraction on the rate of subsequent decision point errors was evaluated. On the basis of the work of Botvinick and Plaut (2004), the number of decision point errors was predicted to be larger following mid-subtask interruption than following end-subtask interruption.

Because we were interested in naturalistic routine action, the experiment was designed to minimize demands that might elicit the application of special mnemonic or control strategies. This necessarily meant that errors in the experiment (as in everyday performance) would be relatively few. In view of this and in order to maximize the sensitivity of error detection, subject performance was evaluated not only for the occurrence of overt, completed errors, but also for what Schwartz, Reed, Montgomery, Palmer, and Mayer (1991) have referred to as "micro-slips," cases in which an erroneous action is initiated but aborted. Our analysis proceeded on the assumption, shared by Schwartz et al., that such partial errors stem from the same factors as completed errors.

METHOD

Subjects

The subjects were 22 University of Pennsylvania students who responded to an electronic newsgroup posting. Eight males and 14 females participated, their ages ranging from 18 to 25 years. The

subjects provided informed consent prior to the beginning of the experiment and were paid by the hour for their participation. In an effort to maximize the yield of errors, the subjects who committed at least one decision point error (defined below) were invited to participate in a second session. Four out of the 9 subjects contacted did return, resulting in a total of 26 subject-sessions.

Procedure

Each session lasted approximately 1 h. During this time, the subject's task was to prepare 50 cups of instant coffee, coping with frequent interruptions to perform a secondary task. The subjects sat at a table on which there was a tray containing a coffee cup and a spoon. Behind the tray stood three large glass canisters, each containing single-serving packets of artificial creamer (left), instant coffee (center), or sugar (right).

At the beginning of each trial, the experimenter placed a cup of hot water on the tray, signaling the subject to begin. Instructions provided at the outset of the experiment established a number of guidelines for performing the coffee-making task. Coffee grounds were always to be added before cream or sugar. Cream and sugar were to be added in a variable order of the subject's choosing, avoiding simple patterns across trials (e.g., alternation). For each ingredient, only one packet was to be added to each cup. Each ingredient was to be stirred in, directly after its addition. Packets were to be thrown away immediately after use. Upon completion of the task, the subject was to place the cup on a counter located to the right of the workspace, where the cup would be retrieved by the experimenter and immediately refilled for the next trial.

Performance of the coffee task was interrupted on 80% of the trials. The timing of this interruption was based on a pseudorandom trial sequence protocol, established before the session and different for each subject. The protocol indicated, for each trial in the session, (1) whether an interruption should take place and, if so, (2) whether it should take place during addition of the second or third ingredient (i.e., during cream or sugar adding, depending on the ordering selected by the subject), and (3) whether the interruption should occur at mid-subtask or end-subtask. Mid-subtask interruptions occurred just following pouring of the relevant ingredient into the cup. To be more precise, the subject was interrupted just as the hand used for pouring began to withdraw from its position above the cup. End-subtask interruptions occurred as the subject completed stirring—that is, just as the hand used for stirring began to withdraw from the position of the cup. The protocol included equal numbers of mid- and end-subtask interruptions, and each of these types of interruption was equally likely to occur during the second and third subtasks.

Interruptions commenced with the experimenter announcing "Freeze!" In response, the subjects were to stop what they were doing, putting down whatever objects they were holding (throwing away packets), and to direct their attention to a 3×5 card, held at eye level by the experimenter. This card presented a change-making problem (a subtraction task involving dollar amounts between \$1.00 and \$9.99). The subjects were asked to state their answer to the problem aloud and then to write it down on a sheet of paper, positioned to the right of the tray. Immediately after this, they were to return to the coffee-making task, picking up where they had left off. All the subtraction problems were of roughly similar difficulty, requiring the carrying of one digit.

At the outset of the session, subjects completed 4 practice trials (2 uninterrupted), after which portions of the instructions were repeated as necessary. Throughout the experiment, reminders of the instructions were provided if the subjects failed to stir in all the ingredients, failed to discontinue the coffee task when interrupted, or added sugar before cream or cream before sugar on more than three consecutive trials. The subjects were offered a short break following the first 25 trials. The entire session was videotaped using a digital camcorder, mounted directly above the workspace.

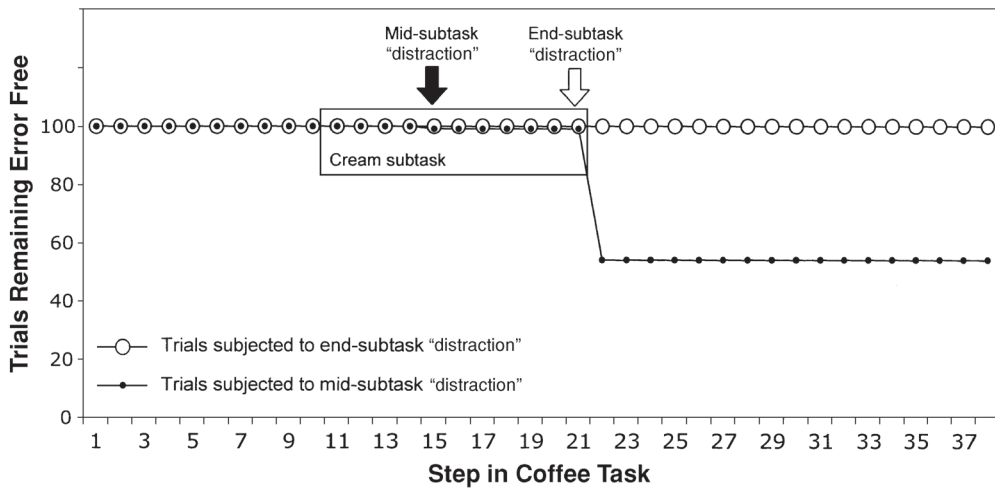


Figure 3. Simulation of momentary distraction, as reported by Botvinick and Plaut (2004). The *x*-axis indicates steps in the coffee-making task. Steps in the cream-adding subtask are boxed. Data points show the number of trials, out of a total of 100, that remained error free at the corresponding step. Arrows indicate the step on which noise was added to the model's context representation, simulating the effect of a brief distraction. Filled points and arrow, mid-subtask interruption; open points and arrow, end-subtask interruption. From "Doing Without Schema Hierarchies: A Recurrent Connectionist Approach to Normal and Impaired Routine Sequential Action," by M. Botvinick and D. C. Plaut, 2004, *Psychological Review*, 111, p. 413. Copyright 2004 by the American Psychological Association. Adapted with permission.

Analysis

Videotapes were viewed independently by two coders, one of whom was not informed of the predictions being tested. Trials were identified that involved either imprecise timing of an interruption or failure on the part of the subject to immediately discontinue the coffee-making task upon interruption. Such trials were excluded from later analyses. The remaining trials were evaluated for the occurrence of *decision point errors*. These were errors involving an incorrect transition from one subtask to the next, as reflected in the execution of an action initiating a context-inappropriate subtask. Given this definition, four specific decision-point errors were possible: (1) reaching toward the sugar canister when sugar had already been added (sugar repetition), (2) reaching toward the creamer canister when creamer had already been added (creamer repetition), (3) reaching for the cup, as if to return it to the experimenter, without having added sugar (sugar omission), and (4) reaching for the cup without having added creamer (creamer omission). These initiating actions were often followed by the remainder of the subtask to which they belonged. However, in some cases, a reach initiated toward an incorrect target was aborted prior to completion, with the hand being retracted and redirected toward a different target. Such partial reaches were counted as decision point errors but were further labeled as *microslips* (following Schwartz et al., 1991). Trials not containing a decision point error were coded as correct.

Two statistical comparisons were conducted. In order to establish whether interruptions led to an increase in errors, subject-sessions were classified according to whether the interrupted trials contained a higher, lower, or equal proportion of errors, as compared with uninterrupted trials, and a sign test was conducted on this data. Second, in order to test the central prediction of the study, rates of decision point errors were compared between trials involving mid-subtask and end-subtask interruptions. Sessions that included at least one decision point error were coded to indicate which interruption condition (mid- vs. end-subtask) produced the higher error rate. A sign test was then performed to evaluate whether one condition yielded more errors, across sessions, than did the other.

RESULTS

Interrater reliability was extremely high. A discrepancy occurred on only a single trial and was resolved by consensus. An average of 2.7 trials per session were excluded from analysis due to experimenter error or failure on the subject's part to discontinue performance immediately upon interruption. The data indicated that the subjects had complied with the instruction to vary the order in which ingredients were added. Trials were coded according to whether sugar was added before cream or cream before sugar, and the balance between these was quantified in terms of informational entropy. Thus, a result of 1.00 would indicate exactly equal frequencies for both orderings. Average entropy across sessions was .99. In order to check whether the subjects adopted a consistent pattern of alternating or repeating orderings, the trials were coded to indicate whether they repeated or reversed the ordering from the previous trial, and this binary coding was used to compute an entropy value for each session, indicating the degree to which the subjects balanced between repeating and alternating orderings. The resulting mean entropy was .92. This submaximal value reflected a bias toward reversing the ordering from the previous trial; across subjects, the mean proportion of such alternation trials was .63, significantly different from the unbiased value of .5 (*t* test, $p < .0001$). This bias toward alternation may have been induced by the instruction to avoid multiple successive order repetitions.

As was anticipated, overall performance was highly accurate, with 1.9% of the trials containing a decision

point error.² Table 1 shows the numbers of trials with and without decision point errors in each interruption condition. Such errors, which took the form of sugar repetitions, sugar omissions, and creamer repetitions, occurred in 50% (13 of 26) of subject-sessions. In every case, a greater proportion of errors occurred on interrupted trials than on noninterrupted trials (sign test, $p = .0002$). Indeed, as Table 1 indicates, no decision point errors at all occurred on noninterrupted trials. As was predicted, decision point errors were observed to be more frequent following mid-subtask interruption than following end-subtask interruption. In the experiment overall, mid-subtask interruptions yielded an error rate of 4.0%, whereas end-subtask interruptions yielded a rate of 0.9%. Of the 13 subject-sessions that contained at least one error, 11 involved a higher proportion of errors in the mid-subtask condition than in the end-subtask condition (sign test, $p = .039$). Similar results were obtained when the same analysis was conducted at the level of subjects, as opposed to sessions, in order to acknowledge that 4 subjects returned to participate in a second session. Of the 22 subjects, 7 (32%) committed at least one error, and for all of these, the proportion of errors was higher following mid- than following end-subtask interruption (sign test, $p = .016$).³

It was noted that the subjects occasionally neglected to stir in an ingredient (usually following a mid-subtask interruption) or inappropriately repeated the stirring procedure (usually after an end-subtask interruption). Such errors were considered potentially problematic, mainly because their occurrence might affect the probability of subsequent decision point errors. In order to exclude this as an explanation for the results of our initial analysis, video data were reviewed for the 13 sessions that contained at least one decision point error, and trials involving a stirring error were identified. Table 1 (bottom) shows the number of correct and decision point error trials after removal of the trials containing a stirring error. Even in this data subset, mid-subtask interruptions were associated with significantly more decision point errors

than were end-subtask interruptions (4.9% vs. 0.5%). Excluding decision point errors that were also accompanied by repetitions or omissions of stirring reduced the count of sessions containing errors to 9. Of these sessions, all 9 involved more errors in the mid-subtask condition than in the end-subtask condition (sign test, $p = .002$).

DISCUSSION

The present experiment tested a novel prediction concerning the effect of distraction on performance in everyday tasks. As was predicted on the basis of computational work by Botvinick and Plaut (2004), the frequency of action slips and, in particular, decision point errors depended on where the distraction occurred relative to subtask boundaries. A greater number of decision point errors followed interruptions occurring midway through a subtask than following interruptions falling at the end of a subtask.

Naturally, further experimentation is called for to establish that the pattern observed in the present study can be reliably replicated and that it generalizes to other tasks. However, if one assumes that this can be shown, the effect appears to provide a highly constraining benchmark for future models both of action slips and of routine sequential action in general.

With respect to slips, the present findings conflict with the idea that errors at decision points are fundamentally due to a failure to execute an attentional check, as has been suggested by earlier theories. Of course, it would be unreasonable to suggest that checking or monitoring operations play no role in naturalistic action. Indeed, the microslips observed in the present experiment may very well reflect the intervention of executive functions, such as those posited by Norman (1981) and Reason (1990, 1992). Nevertheless, the present findings are most consistent with an account that portrays decision point errors as stemming originally not from executive failure, but from a degradation of context information, as has been proposed by Botvinick and Plaut (2003, 2004).

In addition to what the present findings may reveal about action slips, they may also have implications for theories of routine sequential behavior at large. In particular, they lend support to two ideas concerning the way that temporal or task context is represented during the performance of sequential routines. First, the finding that mid-subtask distraction can lead to end-subtask (decision point) errors is consistent with the idea that context information is represented in an active, online fashion, making it susceptible to disruption at any point in the course of performance. Second, the finding that mid-subtask distraction leads to more errors than does end-subtask distraction is consistent with the idea that distinctions between different contexts are attenuated during the performance of subsequences where such distinctions are not immediately relevant to action selection.

Existing models of routine sequential behavior differ in their accounts of how context information is represented, updated, and accessed (see, e.g., Cooper & Shallice,

Table 1
Numbers of Correct Trials and Trials Containing Decision Point Errors for Uninterrupted Trials and Trials Involving Mid- and End-Subtask Interruptions

	Interruption		
	None	Mid-Subtask	End-Subtask
Entire study			
Correct	270	473	467
Decision point error	0	19 (4)	4 (1)
Data subset			
Correct	139	203	201
Decision point error	0	10 (4)	1 (0)

Note—Data are pooled across second- and third-subtask interruption conditions. The upper portion of the table shows the data for the entire study. The lower portion shows the data based on the 13 sessions that contained at least one decision point error, with trials that contained a stirring error excluded. Error counts outside parentheses include both completed errors and microslips. Numbers of microslips are indicated in parentheses.

2000; Grafman, 1995; Miller, Galanter, & Pribram, 1960; Schank & Abelson, 1977). To the extent that the present pattern of findings can be shown to apply generally, it may offer a way of adjudicating among competing accounts. Although a detailed assessment of existing models is beyond the scope of the present article, the present findings do appear to pose a challenge for several well-known accounts. One particularly relevant example is provided by Cooper and Shallice. This study presented a model of routine sequential behavior (contention scheduling), applied to the same task as that addressed by the present experimental study. Within the Cooper and Shallice model, activation of units representing subtasks is gated on the basis of whether the preconditions of the subtask have been met and whether the subtask's goals have been accomplished. Noise in the system can lead to an overriding of this gating mechanism, giving rise to decision point errors, such as subtask omissions and repetitions. However, it is not clear that within the Cooper and Shallice model, a disruption occurring midway through a subtask would lead to an increased number of errors at the transition to the next subtask. This effect might be modeled by allowing noise to affect the gating mechanism's record of accomplished goals. However, even assuming this extension of the model, further changes would be required to account for the differential effect of mid- versus end-subtask distraction.

The present findings also appear to pose a challenge for so-called competitive-queuing models of sequence production (Houghton, 1990), which have been applied to routine sequential behavior and action slips in some recent work (e.g., Humphreys & Forde, 1998). Here again, it is not clear why mid-subtask distraction should lead to an increased number of errors at a subsequent decision point, let alone why distraction at mid-subtask would be more disruptive than end-subtask distraction.

One limitation of the present work stems from the use of interruption as a method of inducing distraction. As has been noted elsewhere (Altmann & Trafton, 2002), the processes involved in coping with interruptions in sequential routines may be rather complex. We assumed that the effect of interruption could be reduced, for present purposes, to nonspecific interference with the representation of the primary task context. If this assumption is valid, it should be possible to produce the same pattern of results using different means to induce distraction. Indeed, analogous findings should result from direct disruption of relevant neural representations of context, induced, for example, by transcranial magnetic stimulation to the dorsolateral prefrontal or the left inferior parietal cortex, regions that have been implicated in disorders of everyday sequential action (Buxbaum, Schwartz, & Montgomery, 1998). In this regard and others, the findings reported here open up interesting opportunities for further investigation.

REFERENCES

ALTMANN, E. M., & TRAFTON, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26*, 39-83.

- BAARS, B. J. (1992). *Experimental slips and human error: Exploring the architecture of volition*. New York: Plenum.
- BOTVINICK, M., & PLAUT, D. C. (2003). Representing task context: Proposals based on a connectionist model of action. *Psychological Research*, *66*, 298-311.
- BOTVINICK, M., & PLAUT, D. C. (2004). Doing without schema hierarchies: A recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, *111*, 395-429.
- BUXBAUM, L. J., SCHWARTZ, M. F., & MONTGOMERY, M. W. (1998). Ideational apraxia and naturalistic action. *Cognitive Neuropsychology*, *15*, 617-643.
- COOPER, R., & SHALLICE, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, *17*, 297-338.
- DUFFY, R. B., & SAULL, J. W. (2003). Errors in technological systems. *Human Factors & Ergonomics in Manufacturing*, *13*, 279-291.
- GRAFMAN, J. (1995). Similarities and differences among current models of prefrontal cortical function. In J. Grafman, K. J. Holyoak, & F. Boller (Eds.), *Structure and functions of the human prefrontal cortex* (Annals of the New York Academy of Sciences, Vol. 769, pp. 337-368). New York: New York Academy of Sciences.
- HECKHAUSEN, H., & BECKMANN, J. (1990). Intentional action and action slips. *Psychological Review*, *97*, 36-48.
- HOUGHTON, G. (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current research in natural language generation* (pp. 287-319). San Diego: Academic Press.
- HUMPHREYS, G. W., & FORDE, E. M. E. (1998). Disordered action schema and action disorganisation syndrome. *Cognitive Neuropsychology*, *15*, 771-811.
- JAMES, W. (1890). *The principles of psychology*. New York: Holt.
- MILLER, G. A., GALANTER, E., & PRIBRAM, K. H. (1960). *Plans and the structure of behavior*. New York: Holt, Rinehart & Winston.
- NORMAN, D. A. (1981). Categorization of action slips. *Psychological Review*, *88*, 1-15.
- REASON, J. T. (1990). *Human error*. Cambridge: Cambridge University Press.
- REASON, J. T. (1992). Cognitive underspecification: Its varieties and consequences. In B. J. Baars (Ed.), *Experimental slips and human error: Exploring the architecture of volition* (pp. 71-91). New York: Plenum.
- REASON, J. T., & MYCIELSKA, K. (1982). *Absent minded? The psychology of mental lapses and everyday errors*. Englewood Cliffs, NJ: Prentice Hall.
- SCHANK, R. C., & ABELSON, R. P. (1977). *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Erlbaum.
- SCHWARTZ, M. F. (1995). Re-examining the role of executive functions in routine action production. In J. Grafman, K. J. Holyoak, & F. Boller (Eds.), *Structure and functions of the human prefrontal cortex* (Annals of the New York Academy of Sciences, Vol. 769, pp. 321-335). New York: New York Academy of Sciences.
- SCHWARTZ, M. F., REED, E. S., MONTGOMERY, M., PALMER, C., & MAYER, N. H. (1991). The quantitative description of action disorganisation after brain damage: A case study. *Cognitive Neuropsychology*, *8*, 381-414.
- ZACKS, J. M., & TVERSKY, B. (2001). Event structure in perception and conception. *Psychological Bulletin*, *127*, 3-21.

NOTES

1. In the present article, we use the term *subtask* informally and intuitively to refer to the sequences of actions that come between decision points. The question of what defines a subtask is an interesting and complex one, which we have addressed in detail elsewhere (Botvinick & Plaut 2003, 2004; see also Zacks & Tversky, 2001). However, for present purposes, the only property required of the sequences occurring between decision points is that they be insensitive to the context information that is brought to bear at the decision point that marks their end. Thus, for example, in coffee making, the sequence used for adding cream takes essentially the same form regardless of whether sugar has yet been added, although this information *does* become relevant at the end of the cream-adding sequence.

2. The mean trial number for errors was 19. There was a tendency for errors to occur in the first half of the experimental session: Seventeen errors occurred in the first half of the session, 6 in the second. The number of errors was higher in the first half of the experiment for 10 sessions, and in the second half for 1 session (sign test, $p = .012$).

3. Although it is not one of the questions our experiment was designed to address, it is interesting to consider whether there were differences between the sessions in which errors occurred and the sessions in which they did not. There was no apparent difference between these two groups of sessions in terms of biases toward one ordering of ingredients; entropy (as introduced earlier) was .99 in both cases. However, a difference was noted when entropy was computed on the basis of order repetitions versus alternations: For the sessions in which no error occurred, this number was .96; for the sessions with errors, it was .89 (t test, $p = .031$).

Further investigation revealed that the subjects were 6% more likely to alternate ingredient order between trials in the sessions in which errors occurred than in the error-free sessions (66% vs. 60%), although this difference did not quite reach statistical significance (t test, $p = .07$). More errors occurred on alternation trials—that is, trials that reversed the order of the cream and sugar subtasks from that in the preceding trial (16 vs. 6). However, this was not significantly different from what would be expected by chance, given the greater frequency of alternation trials overall, as indicated by a sign test comparing the number of sessions with higher error rates on alternation than on repetition trials with the number of sessions showing the reverse pattern ($p = .27$).

(Manuscript received March 23, 2004;
revision accepted for publication March 24, 2005.)