

How Long Have I Got? Making Optimal Visit Durations in a Dual-Task Setting

George D. Farmer, Christian P. Janssen, and Duncan P. Brumby

george.farmer@gmail.com, c.janssen@ucl.ac.uk, d.brumby@cs.ucl.ac.uk

UCL Interaction Centre, University College London, Gower Street, London, WC1E 6BT, UK

Abstract

Can people multitask optimally? We use a dual-task paradigm in which participants had to enter digits while monitoring a randomly moving cursor. Participants earned points for entering digits correctly and were docked points if they let the cursor drift outside of a target area. The severity of the tracking penalty was varied between conditions. Participants therefore had to decide how long to leave the tracking task unattended. As expected, participants left the tracking task for longer when the penalty was less severe and also when the cursor moved less erratically. To test whether participants were adjusting their behavior in an optimal manner, observed behavior was compared to a prediction of the optimal visit duration for each condition. Overall, the degree of correspondence between the observed behavior and the predicted optimum was very good, suggesting that people can multitask in a near optimal fashion given explicit feedback on their performance.

Keywords: Multitasking; Objective functions; Optimal Performance.

Introduction

In many everyday settings people are faced with the problem of deciding how much time to spend on a given task or activity. For instance, should you take the time to read this paper? This decision becomes more difficult to make if you have several other tasks or activities that also need to be completed. There has been a great deal of interest recently in how people coordinate and manage multiple tasks and activities (e.g., Payne, Dugan, & Neth, 2007; Salvucci & Taatgen, 2011). Understanding the limits on our abilities to multitask well has important theoretical and practical implications. In this paper, we focus on a core aspect of multitasking behavior: deciding how much time to spend on a given task before switching to another. Are people any good at making this decision in a dual-task setting, given clear feedback on their performance?

To address this question we use an explicit payoff regime to objectively define what good dual-task performance is in the context of our study. Payoff regimes have been used in a number of studies (e.g., Janssen et al., 2011; Schumacher et al., 1999; Trommershäuser, Maloney, & Landy, 2008). One of the benefits of using a payoff regime is that it allows multiple task performance metrics to be combined into a single currency (or value) that can be communicated to the participant. For instance, in a simple stimulus-response task, participants might be given more points for making faster responses but docked points for making an error. The participants' pay might then be determined by the points accrued during the study.

A major benefit of using a payoff regime in an experiment is that it provides an objective criterion against which optimal task performance can be defined: People should select behaviors that maximize reward (Howes, et al., 2009). The idea that people adapt behavior to maximize reward is a key element of a number of important theoretical approaches to understanding human behavior, including models based on decision theory in economics, reinforcement learning, signal detection theory, and some Bayesian approaches.

A recent study by Janssen, Brumby, Dowell, Chater and Howes (2011) gives support to the idea that people can optimally adapt their behavior in a complex dual-task setting to maximize reward. In their study, participants had to type a string of 20 digits while keeping a moving cursor inside a target area. Participants were given explicit feedback on their performance at the end of every trial using a payoff function that integrated performance across the two tasks into a single value. The payoff function rewarded participants for entering the digit string quickly but penalized them for allowing the cursor in the tracking task to move outside of its target area. A model was developed to identify the task interleaving strategy that would maximize payoff for different conditions, which varied how difficult it was to keep the cursor inside the target area. Janssen et al. (2011) showed that participants' observed strategies were very near the optimal strategy for each condition.

A limitation of the Janssen et al.'s (2011) study was that only a single payoff regime was used. A stronger test of the hypothesis that people select behaviors that maximize reward (e.g., Howes, et al., 2009) would be given by varying the payoff regime and seeing whether people adapt their behavior accordingly.

In this paper, we use a dual-task paradigm and vary the nature of the payoff function used to reward participants. We use a modified version of the dual-task paradigm developed by Janssen et al (2011), in which participants had to enter digits while keeping a moving cursor inside a target area. As in Janssen et al.'s study, participants were rewarded for entering digits correctly and were penalized for allowing the cursor to drift outside of its target area. The severity of this *tracking penalty* was varied between conditions, such that participants either lost all the points gained after entering a set of digits, half of the points gained, or incurred a fixed penalty. Given this change in tracking penalty, we were interested in whether participants would adjust their task interleaving strategy in such a way as to maximize the payoff achieved in each condition.

The current study was designed in such a way to give participants considerable flexibility in deciding how to allocate their time across the two tasks over an extended period of time. In contrast, participants in Janssen et al.'s (2011) study were required to enter 20 digits in the typing task. Participants could do this fairly quickly, and the fixed number of digits meant that there was a limited space of plausible ways to interleave the two tasks. The current study had participants enter a continuous stream of digits while monitoring the tracking task over a period of 120 seconds. This meant that participants had considerably greater freedom in deciding how long they should leave the tracking task unattended for while entering digits.

This change in the structure of the dual-task setup from that used in Janssen et al. (2011) meant that participants in the current study could also be given feedback on their performance, in terms of the number of points earned, after every visit to the typing task as well as at the end of every trial. We were primarily interested in whether participants would adjust their behavior in such a way as to maximize the cumulative payoff achieved over a trial in each condition. To do this, we compare observed performance to a prediction of the optimal visit duration for each condition, which was inferred using a model of the task environment and observations of basic task execution times.

Before describing the experiment and the results in detail, we first briefly consider why participants might not be able to perform optimally in this dual-task setting. In our study, participants will receive feedback on their performance both at the end of every visit to the typing task and at the end of the trial. That is, they will be exposed to local feedback on their immediate behavior and global feedback on their overall performance. The participant's task is to achieve the best possible score at the end of the trial based on the incremental accumulation of points during the trial. This is a fairly complex assessment to make because the number of visits that can be made in a trial is not fixed but is instead dependent on the duration of each visit: The participant must choose between fewer longer visits and many shorter visits. Moreover, participants were not given explicit information on the underlying form of the payoff function, but instead had to evaluate their behavior after receiving feedback. In this context, participants might meliorate by adopting a strategy that maximizes the highest local reward (e.g., the reward value after a single visit) rather than finding the strategy that maximizes the cumulative reward to be had at the end of the trial (e.g., Gureckis & Love, 2009; Neth, Sims, & Gray, 2006).

Method

Participants

Twenty Master's students (seven female) from University College London participated on a voluntary basis. Participants were aged between 22 and 37 years ($M = 27$ years). An incentive of a £10 voucher was offered for the participant who achieved the highest score in the study.

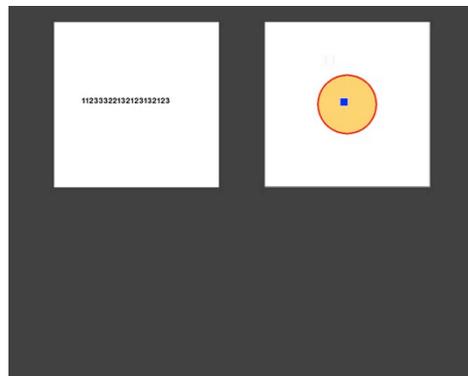


Figure 1: Position of the two task windows in the interface. Only one of the task windows was visible at a time.

Materials

A dual-task setup similar to that of Janssen et al. (2011) was used, in which participants completed a discrete typing task while monitoring a continuous tracking task (see Fig. 1). Tasks were displayed on a 17-inch monitor at a resolution of 1024 by 1280 pixels, with each task being presented within a 450 x 450 pixel area. The typing task was presented on the left of the monitor and participants entered digits with their left-hand using a numeric keypad. The tracking task was presented on the right of the monitor and participants controlled the task with their right-hand using a Logitech Extreme 3D Pro joystick. There was a horizontal separation of 127 pixels between the two tasks. At any moment in time only one of the tasks was visible on the monitor. By default the typing task was visible and the tracking task was covered by a gray square. Holding down the trigger of the joystick revealed the tracking task and covered the typing task. Releasing the trigger covered the tracking task and made the typing task visible again. Each task could only be controlled when it was visible.

The typing task required participants to enter digits using a numeric keypad. A continuous list of to-be-entered digits was generated from the numbers 1 to 3 drawn in a random order with the constraint that no digit was repeated more than three times in a row. At any moment, 27 digits were shown on the typing task display. As a participant typed, the left-most digit in the list would disappear, all digits would then move one to the left, and a new digit would appear in the right-most position. The display would remain unchanged if an incorrect digit were entered, meaning that the participant did not have to rectify the error. However, a *typing penalty* was applied through the payoff function.

For the tracking task participants were required to keep a square cursor inside a circular target area. The cursor was 10 x 10 pixels and the target area had a radius of 120 pixels. The movement of the cursor was updated every 23 milliseconds. Values were sampled from a Gaussian distribution to determine the size of the cursor's movement, and the parameters on this distribution were varied between conditions in such a way to make the cursor move about at different speeds. Holding down the trigger of the joystick allowed participants to see the tracking task and move the

cursor around with the joystick if they so wished. Releasing the trigger covered up the tracking task again.

Due to the nature of the random drift function, the cursor could move outside of the target area, and move back in again, whilst the participant was typing. This would make it difficult for the participant to know whether the cursor had moved outside the target area while they had been working on the typing task. We therefore used a change in cursor color to make it clear whether or not the cursor had moved outside the target area. At the start of the trial the cursor was blue. If it moved outside of the target area during a typing visit it changed to red. The cursor returned to blue once the participant returned the cursor inside the target area.

Using this dual-task setup, participants completed a series of trials, each lasting for 120 seconds. During each trial the main decision facing the participant was to judge how long they should leave the tracking task to enter digits. Participants received feedback on their performance in terms of number of points achieved under the payoff function after each visit to the typing task. The payoff function rewarded participants for each digit that was correctly entered during the visit, and penalized them for entering digits incorrectly. An additional penalty was also applied if the cursor drifted outside of the target area during the visit. The precise nature of the payoff function was varied between conditions. Feedback was displayed above the tracking task and remained visible while the tracking task was being worked on. This allowed the participant to evaluate their performance after each visit to the typing task. Cumulative feedback was also given at the end of each trial.

Design

A 3x2 (payoff function x cursor noise) within-subjects design was used. After every visit to the typing task participants received feedback on their performance in terms of the number of points achieved. Participants received 10 points for every digit that was typed correctly and were docked 5 points for every digit that was typed incorrectly. Points were also deducted if the cursor drifted outside of the target area. The severity of this tracking penalty was varied between conditions: Participants either lost all the points gained for that visit (lose-all condition), half of the points gained for that visit (lose-half condition), or incurred a fixed penalty of 500 points (lose-500 condition).

The study also implicitly varied the amount of time that the tracking task could be left unattended. Recall that the movement of the cursor was updated every 23 milliseconds. Movement values were sampled from a Gaussian function with a mean of zero and a standard deviation of either three (low noise condition) or five (high noise condition) pixels per update. This meant that in the high noise condition, the cursor moved more erratically and was more likely to travel outside of the target area sooner than in the low noise condition (see Fig. 2).

The dependent variable of interest was the mean visit duration to the typing task over each trial. This measure captures the tradeoff that participants had to make between

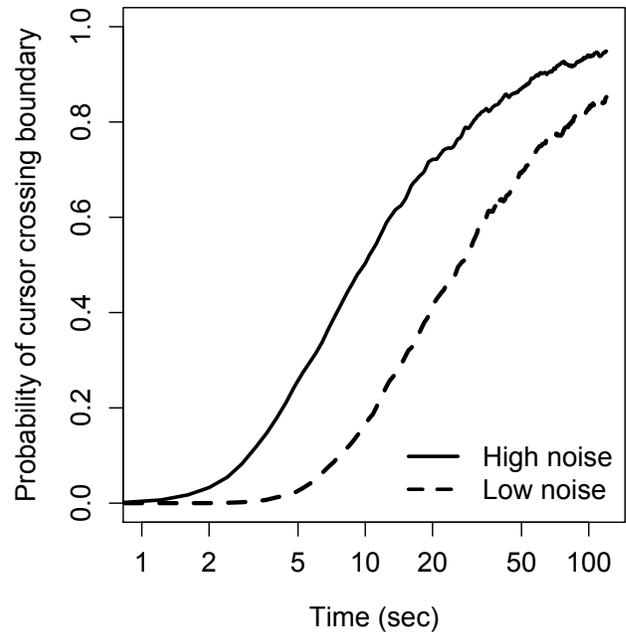


Figure 2: Probability of cursor crossing the target area boundary if left unattended during the course of a trial in the high and low noise condition.

gaining points (by typing more digits), and losing points (by incurring the tracking penalty) in this dual-task setting.

Procedure

Upon arrival, participants received instructions on the dual-task setup. Crucially, they were briefed that they could gain points in dual-task trials through fast and accurate typing and that they would lose points by making typing errors and by letting the cursor drift outside the target area. Participants were not informed of the exact way in which the gain and penalties were calculated in each condition. They were told that the payoff function changed between blocks of trials.

At the beginning of the experiment participants were given a chance to practice each of the tasks separately (tracking for two trials of 10 seconds each and typing for two trials of 20 seconds each) before being given a chance to practice the tasks together (for two trials of 30 seconds each). Following the practice session, participants completed six experimental blocks (one for each condition). For each block, participants first completed two single-task tracking trials of 10 seconds each so that they could estimate the speed of the cursor, and two single-task typing trials of 20 seconds each. There were then two dual-task trials of 120 seconds each. We report data from only the second of each dual-task trial from each block (condition). In total the experiment took about 45 minutes to complete. Participants were given a brief break after every other block of trials. The order in which the different payoff conditions were experienced was randomized and counter-balanced across different participants. The order of the noise conditions was assigned randomly for the first two blocks, but repeated for each payoff condition.

Model

Before describing the results of the study, we first describe the model that was used to derive predictions for optimal performance for each condition. We assume that optimal performance corresponds to behavior that maximizes the payoff achieved over the trial. In order to maximize payoff, participants had to determine how long to make each visit to the typing task before returning to the tracking task. If a visit is too long, then a penalty will likely be incurred (as the probability of the cursor leaving the target increases). If a visit is too short, then the participant will incur needless switch costs (from flicking back and forth between tasks). For each condition we assume that there will be a point of optimal gain. We formalize these intuitions with a model.

We start by formalizing the drift rate of the cursor for the two noise conditions. Figure 2 shows the probability of the cursor leaving the target area during a given period of time for the two noise conditions. These functions were derived from 3,000 simulations in which the cursor was started at the center of the target display and left to drift for 120 seconds. We then calculated the mean distance of the cursor from the center of the target area at each time step across each run. As expected, the cursor was more likely to cross the target area boundary sooner in the high noise condition.

With an estimate of the rate of drift for the tracking task, we develop a simple model of the typing task. In order to develop this model we take the average time it took a participant to enter a digit. We observed that the mean interval between two keypresses was 40 msec ($SD = 10$ msec) across all dual-task conditions. We therefore use this estimate of 40 msec in the model. We also incorporated typing errors into the model by fixing the error rate to that which was observed across all dual-task conditions. It was found that on average 6% of all key presses made during a trial were incorrect. We therefore had the model make extra (incorrect) keypresses at this rate.

After typing a number of digits, the model switched from typing to tracking. We found that the mean visit duration to the tracking task was 1.32 sec ($SD = 0.36$ sec) across the various dual-task conditions. Based on this observation we made the assumption that each visit to the tracking task would always be 1.32 sec (regardless of the actual position of the cursor). We also made the assumption that during each visit to the tracking task, the cursor was brought back to the center of the target area. This meant that the time for the cursor to cross the boundary was the same for each visit.

With the basic components of the model in place we can consider the payoff that would be achieved by making visits of different durations to the typing task. We consider 300 simple strategies in which a consistent number of digits were entered during each visit to the typing task display. The number of digits entered ranged between typing one digit per visit (frequent interleaving) to typing 300 digits in a single visit (i.e., typing the entire trial without tracking). To assess the success of each strategy, we calculated the expected payoff of each strategy for each experimental condition.

Table 1. Mean duration of visits to the typing task (in sec) for human data and predicted optimal durations.

	Low noise			High noise		
	Half	All	500	Half	All	500
Data	5.25	4.63	4.03	3.25	2.47	2.18
Model	8.00	6.00	4.81	6.00	3.61	1.61

Note: For the lose-500 condition the true optimal prediction was for visit durations of 120 seconds. The reported values are for the earlier peak in the payoff curve (see Fig. 3).

Results

We consider the mean duration of visits to the typing task during each 120-second trial. Before presenting the results of the modeling analysis used to identify the optimal visit duration for each condition and whether participants achieved it, we first consider the impact of the experimental manipulations on performance.

Table 1 shows mean visit durations to the typing task for the second dual-task trial for each condition. It can be seen that participants made visits of different durations dependent on the payoff function. On average, visits to the typing task were longest in the lose-half condition and shortest in the lose-500 condition. Cursor noise also had an effect, in that, participants made shorter visits to the typing task in the high noise condition than in the low noise condition. For statistical analysis a 3x2 (payoff function x cursor noise) repeated-measures analysis of variance (ANOVA) was used, finding significant main effects of payoff function, $F(2,38)=5.1$, $p<.05$, and cursor noise, $F(1,19)=61.12$, $p<.001$. The interaction was not significant, however, $F<1$. In short, participants were adjusting their behavior to the different experimental conditions. We next use the results of the modeling analysis to determine whether the observed visit durations were optimal.

Figure 3 shows a series of data plots showing the duration of visit to the typing task against the payoff achieved for each condition. In each data plot the human data is shown along with the predicted *cumulative* payoff that would be achieved by making visits of a given duration over the 120-second period of the trial (the dark grey lines).

To derive these predictions of the cumulative payoff for each condition, we first calculated the expected payoff for a single visit of a given duration (shown as the light grey lines in the figure). Recall that we systematically varied the number of digits that the model entered during each visit to the typing task. Given that we assumed that entering a digit would take 40 msec on average, it was possible to infer how long a visit would be for a strategy that entered a given number of digits. Given this estimate of visit duration, it was then possible to determine the probability that the cursor would leave the target area during that visit (see Fig. 2). In this way, it was possible to calculate the average payoff that would be expected for a given visit duration. These values are shown in Fig. 3, which shows the payoff that would be expected after a single visit to the

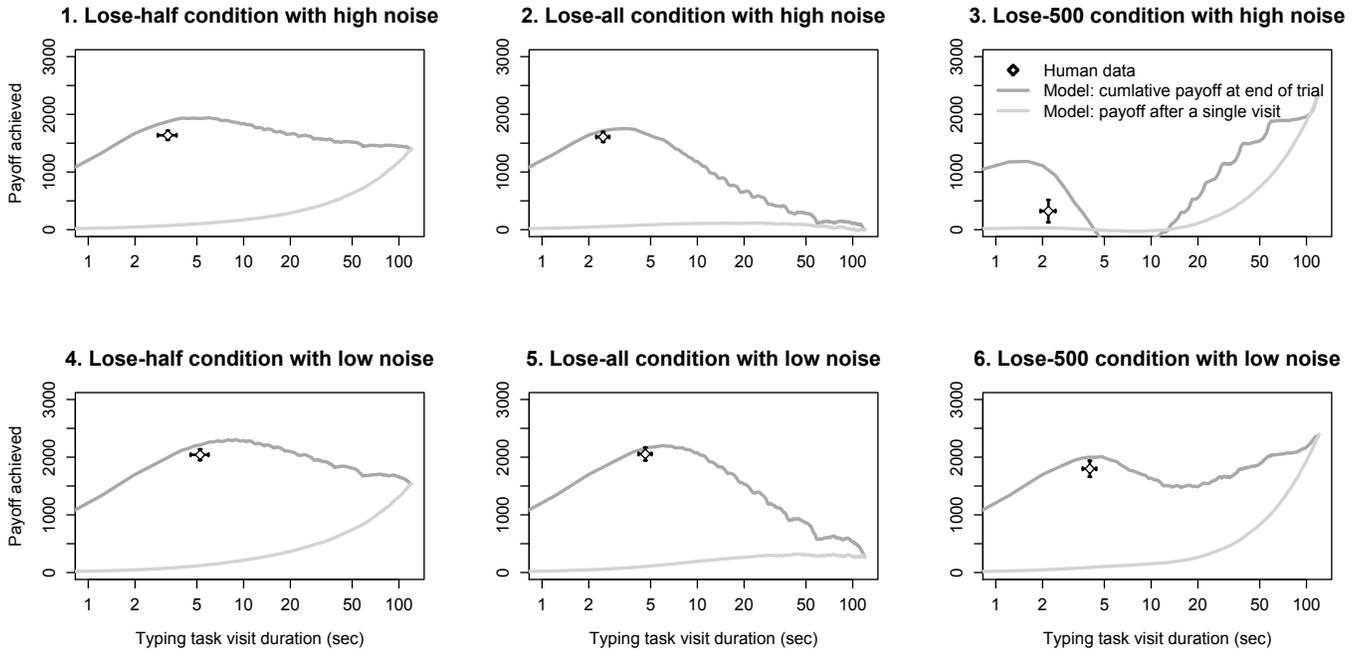


Figure 3: Data plots showing payoff achieved against duration of visit to the typing task for each condition. Diamond data point represents mean human performance (error bars represent standard error of the mean). Model curves show the payoff achieved after a single visit to the typing task (light grey line) and cumulative payoff at the end of the trial (dark grey line).

typing task. To calculate the cumulative payoff achieved for a given strategy, we had to incorporate the time given to the tracking task between visits. As outlined above, we assumed that all visits to the tracking task were 1.32 seconds in duration. With this time cost added, we calculate the number of visits of a given duration that could be made within the 120-second period of the trial and the cumulative payoff that would be expected.

It can be seen in Fig. 3 that the optimal visit duration (i.e., the peak of the curve) is different for each condition. The payoff curves for the lose-half and lose-all conditions have a single defined peak. The reason for this is fairly intuitive: Very short visits earn few points and needlessly give too much time to tracking, while very long visits are likely to incur the penalty and lose points. Furthermore, it can be seen in the figure that the optimal visit duration varies between conditions: Visits should be shorter in the lose-all condition than in the lose-half condition, and visits should be shorter in the high noise condition than in the low noise condition.

The lose-500 condition provides an interesting exception to the above pattern. Here the payoff curve has two defined peaks. The most points are to be gained by making a single very long visit to the typing task. This is because the penalty is fixed and the loss can be made up over the course of the trial by entering digits. There is however a second, lower peak to be had in the lose-500 condition by making many very short visits to the tracking task.

Comparing the human data to the cumulative payoff curves it can be seen that in the lose-all and lose-half payoff conditions, participants were making visits of a duration that

were remarkably close to the peaks of the payoff curves. For the lose-500 condition, not one single participant adopted the optimal strategy, which was to ignore the tracking task completely and make a single visit to the typing task. Instead, participants honed in on the earlier peak in the payoff curve by making very short visits.

For statistical analysis, we compare human data to the predicted optimal visit duration time (i.e., the peak of the curves in Fig. 3). Table 1 summarizes these data. Note that because the lose-500 condition had two peaks, we compared the human data with the earlier peak since no single participant adopted the extreme strategy of completely ignoring the tracking task. In general, it can be seen that there is a very good degree of correspondence between the observed visit durations for each condition and those predicted to be optimal by the model, $R^2 = 0.79$, $RMSE = 1.79$ sec. However, the RMSE is relatively high. This is because participants made shorter visits on average in the lose-half condition than the predicted optimal duration.

Discussion

We investigated whether people can hone their behavior in a dynamic dual-task setting to maximize reward from a payoff function. Participants had to determine how long they could leave a dynamic tracking task unattended to enter digits in a continuous typing task. The payoff function rewarded participants for entering digits correctly but penalized them for allowing the cursor to drift outside of the target area. By manipulating the penalty function, the optimal time to make each visit to the typing task differed between conditions.

Results showed that participants adjusted to the change in penalty function by making longer or shorter visits according to what would maximize the cumulative payoff at the end of the trial. That is, participants honed in on the optimal strategy for each condition. What is remarkable about this is that participants were not simply learning to maximize the amount of points gained after a single visit. If they were doing this, then they would have made very long visits in most of the conditions (see Fig. 3). Instead participants were integrating feedback on local rewards with a fairly accurate assessment of how many visits they would be able to fit within the trial so as to maximize the cumulative payoff at the end of the trial. In other words, they were assessing the rate of reward. These results are consistent with the general idea that people select behaviors that maximize reward (e.g., Howes, Lewis & Vera, 2009; Janssen et al., 2011; Trommershäuser et al., 2008).

The optimal visit duration for each condition was inferred using a model. It should be stressed that these are not model fits in the traditional sense: the model predictions are derived from systematically exploring the space of possible strategies and making a prediction based on the performance characteristics of the strategy that maximizes payoff (see, Howes, Lewis, & Vera, 2009). With the exception of the lose-500 condition, the degree of overall correspondence between the observed visit duration and the optimal model was very good.

In the lose-500 condition participants did not adopt a strategy that achieved the maximum payoff. The reason for this is that the cumulative payoff curve in this condition had two peaks. The most points were to be gained in this condition by making a single very long visit to the typing task. This extreme strategy would have required the participants to completely ignore one of the two tasks. It seems reasonable to think that participants would not spontaneously adopt this behavior as they were told that they were taking part in a *multitasking* study. Moreover, participants received feedback when they switched between tasks. Therefore, making many shorter visits to the typing task is far more informative as feedback is given more frequently, especially since participants were given quite limited exposure to the environment (i.e., two trials per condition).

One issue with the model is that we assumed that all visits to the tracking task were of a constant duration. It seems reasonable to think that time spent on the tracking task might be dependent on the distance of the cursor from the center of the target area. In this modeling analysis we chose not to incorporate this aspect. If we had done it would have given a second dimension on which strategies could vary (see Janssen et al., 2011). For the current model this dimension was probably not as interesting, as most participants spent a relatively fixed amount of time in the tracking window to return the cursor to the middle of the target area. However, there might be more subtle differences that are not being accounted for in the current model. For instance, this constant time assumption seems particularly

problematic for very short visits: as here the cursor might not have moved very far at all and a correction might not be needed at all.

In summary, we have shown that people are sensitive to feedback on performance in a complex dual-task environment. The participants in our study made fairly sophisticated judgments about the amount of time to spend on one task before switching to another, and in that way achieved near optimal performance (in terms of received feedback). What is remarkable to us is that participants were able to do this despite being given very little exposure to the tasks and the payoff regime (i.e., only two trials per condition). Future work might consider what learning processes underlie this ability and how people represent and reason about feedback information.

Acknowledgments

The work reported here was conducted as part of the first author's M.Sc. research project. The other authors were supported by EPSRC grant EP/G043507/1.

References

- Gureckis, T. M., & Love, B. C. (2009). Short-term gains, long-term pains: How cues about state aid learning in dynamic environments. *Cognition, 113*, 293-313.
- Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action. *Psychological Review, 116*, 717-751
- Janssen, C. P., Brumby, D. P., Dowell, J., Chater, N., & Howes, A. (2011). Identifying optimum performance trade-offs using a cognitively bounded rational analysis model of discretionary task interleaving. *Topics in Cognitive Science, 3*, 123-139.
- Neth, H., Sims, C. R., & Gray, W. D. (2006). Melioration dominates maximization: Stable suboptimal performance despite global feedback. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th annual meeting of the cognitive science society* (pp. 627-632). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Payne, S. J., Duggan, G. B., & Neth, H. (2007). Discretionary task interleaving: Heuristics for time allocation in cognitive foraging. *Journal of Experimental Psychology: General, 136*, 370-388.
- Salvucci, D.D. & Taatgen, N.A. (2011). *The Multitasking Mind*. Oxford University Press.
- Schumacher, E.H., Lauber, E.J. Glass, J.M. Zurbriggen, E.L. Gmeindl, L., Kieras, D.E., & Meyer, D.E. (1999). Concurrent response-selection processes in dual-task performance: Evidence for adaptive executive control of task scheduling. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 791-814.
- Trommershäuser, J., Maloney, L.T., Landy, M.S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Science, 12*, 291-297.