Towards optimal payoff manipulations

Christian P. Janssen[1], Andrew Howes[2], and Duncan P. Brumby[1]

[1]UCL Interaction Centre, University College London
London WC1E 6BT UK

[2]School of Computer Science, University of Birmingham
Birmingham, B15 2TT UK

c.janssen@ucl.ac.uk, howesa@bham.ac.uk, d.brumby@cs.ucl.ac.uk

Keywords: Payoff function, mathematical modeling, performance trade-offs, cognitively bounded rational analysis

Introduction

Multitasking typically requires people to make performance trade-offs: paying more attention to one task can improve performance there, but might lead to performance decrements on other unattended tasks. In our work we try to gain a better understanding of how people make such trade-offs. One difficulty in this effort is that performance is typically expressed in different units across tasks (e.g., "accuracy" of keeping a car inside a lane and "speed" of performing a secondary task such as dialing). How do people trade-off these different units?

Explicit payoff functions have been proposed as a way to achieve the desired trade-off (e.g., Howes, Lewis, & Vera, 2009; Janssen, Brumby, Dowell, Chater, & Howes, 2011; Payne, Duggan, & Neth, 2007; Schumacher, et al., 1999). They can be used to translate performance on multiple tasks into a single score. The participant and the modeler can then use this feedback to objectively compare performance for different strategies (Howes, et al., 2009; Janssen, et al., 2011). If successful, payoff functions can be used as a formal way to manipulate a user's priorities. Different strategies can be made optimal through changes of the payoff function. In ongoing work we are exploring under what conditions such "ideal payoff manipulations" can be made. We report some of our intermediate findings here.

Ideal payoff functions and manipulations

A payoff function is a function that translates *performance* on one or multiple tasks into a *single, explicit, objective currency* in a *consistent* manner. As a rule-of-thumb, the output of the function should be meaningful to the participant. Participants can then use the output of the payoff function to assess how well they are performing by comparing payoff values across trials and strategies. Similarly, a model can be used to compare payoff values of different strategies (e.g., Howes, et al., 2009; Janssen, et al., 2011). In this way, a *payoff curve* can be generated that captures how payoff fluctuates as a function of the possible strategies. An ideal payoff curve has four properties:

1. It has one global maximum.
2. It has no local maxima other than the global maximum.
3. The set of strategies that has a payoff value close to the maximum is narrow and these strategies are very similar in nature to the "optimum" strategy.
4. The distribution of possible payoff values is consistent and narrow, such that the mean payoff value of a strategy is representative of the distribution of values.

In essence, these properties guarantee that there is a unique, consistent, clear optimum strategy. This makes it easier for the participant and the modeler to identify the optimum strategy and to assess whether participants performed optimally. The black line in Figure 1 is an example of how payoff score (vertical axis) changes as a function of strategy (horizontal) in an ideal payoff curve.

If payoff is successful in manipulating performance, then in an ideal setting this can be used to make any arbitrary strategy "optimal" for at least one payoff function. In Figure 1 multiple alternative payoff curves are plotted in grey lines, as generated by hypothetical alternative payoff functions. This is considered an *ideal payoff manipulation*, because:

1. Each curve is an ideal payoff curve.
2. Across curves, each strategy is the optimum of at least one ideal payoff curve.

With this definition of an ideal payoff manipulation, we are currently exploring under what conditions (e.g., what types of tasks) such ideal manipulations are possible.

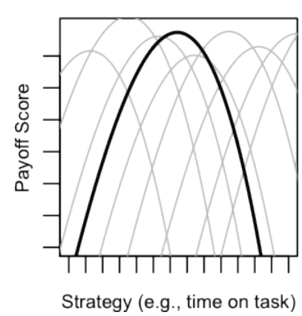

Figure 1: An ideal payoff manipulation (see text)

A mathematical model of interleaving

Inspired by our previous work on a tracking-while-typing task (Farmer, Janssen, & Brumby, 2011; Janssen, et al., 2011), we developed a mathematical model of this task. The model had to keep a one-dimensional first-order moving cursor inside a target area. The movement of the cursor was modeled using Pascal's triangle, which can be used to give exact predictions of the probabilities of the position of the cursor at each timestamp. The model could exert active control of the cursor to overwrite the random movement. In addition a simple secondary task was included, which solely involved opening a task window.

The model could only attend to one task at a time and experienced switch costs when switching between tasks. Both tasks were kept extremely simple on purpose, as this allowed us to focus on the role of payoff functions and whether an ideal payoff manipulation was possible, without having to worry about the correctness of each constraint and about the effects of constraints on performance.

Performance on both tasks was encapsulated in the payoff function. For the tracking task, the model gained points on every sample when the cursor was inside its target area; it lost points otherwise. For the secondary task, the model gained points whenever the window of this task was open; it lost points otherwise. The values of these four gain and loss components were systematically manipulated to explore whether ideal payoff manipulations were possible. We also explored the effects of using a log or exponential transformation to the functions. The general pattern of results was similar across these simulations.

Results and Discussion

As a first step to identify ideal payoff manipulations, we explored whether the location of the global maxima differed across payoff functions. In contrast to the definition of an ideal payoff manipulation, the strategy required to achieve the maximum payoff did not vary much as the payoff function varied. Only specific points emerged as maxima. These maxima were for strategies of which the performance of at least one of the underlying tasks (i.e., how many time units was the cursor inside the target area, how many time units was the secondary task window open) had a local maximum. These local maxima were themselves the result of the constraints imposed by the task environment (e.g., boundary of the tracking task target) and cognition (e.g., switch costs). That is, global maxima in the payoff curve emerged at positions where the interaction of the constraints led to beneficial performance trade-offs.

Looking at individual curves, many also violated the characteristics of an ideal payoff curve. Figure 2 shows three example curves. For each curve (different color lines), the strategy with the highest payoff is highlighted with an open circle. As can be seen, the curves violate the properties of an ideal payoff curve. There are local maxima and there are many strategies that achieved values close to the maximum value. This implies that fine tuned attentional strategies are not always required. We are therefore making further efforts to find task variants in which more subtle strategy choices are required.

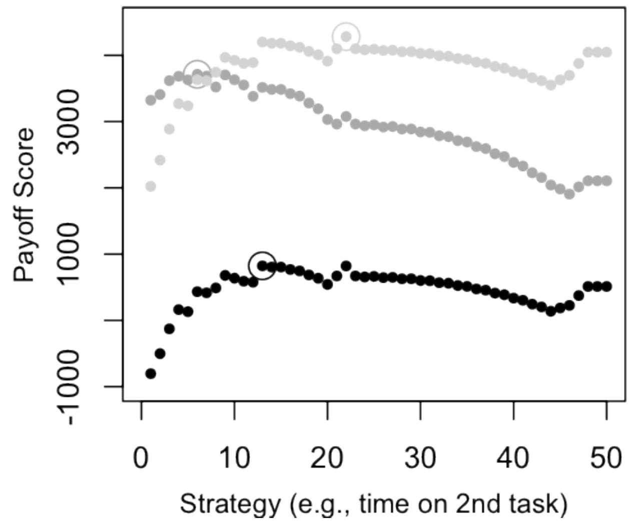

Figure 2: Three example payoff curves

Acknowledgments

This work was sponsored by EPSRC grant EP/G043507/1.

References

Farmer, G. D., Janssen, C. P., & Brumby, D. P. (2011). How long have I got? Making optimal visit durations in a dual-task setting. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society* (pp. 862-867). Austin, TX: Cognitive Science Society.

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.

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.

Schumacher, E., Lauber, E., Glass, J., Zurbriggen, E., 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.