

Received Date: 24-Aug-2012

Revised Date: 20-Mar-2015

Accepted Date: 13-Apr-2015

Article Type: Regular Article

Predicting Short-Term Remembering as Boundedly Optimal Strategy Choice

Andrew Howes^{*}, Geoffrey B. Duggan[†],
Kiran Kalidindi^{**}, Yuan-Chi Tseng^{***}, Richard L. Lewis^{††}

^{*}Department of Computer Science,
University of Birmingham

[†] Department of Psychology,
University of Exeter

^{**}Manchester Business School
University of Manchester

^{***}Department of Industrial Design, National Cheng Kung University

^{††}School of Psychology
University of Michigan

HowesA@bham.ac.uk

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/cogs.12271

This article is protected by copyright. All rights reserved

August 2014

PLEASE DO NOT QUOTE OR DISTRIBUTE

Abstract

It is known that on average people adapt their choice of memory strategy to the subjective utility of interaction. What is not known is whether individual choices are *boundedly optimal*. Two experiments are reported that test the hypothesis that individual decisions about the distribution of remembering between internal and external resources are boundedly optimal where optimality is defined relative to experience, cognitive constraints, and reward. The theory makes predictions that are tested against the data, not fitted to it. The experiments use a no-choice/choice utility learning paradigm where the no-choice phase is used to elicit a profile of each participant's performance across the strategy space and the choice phase is used to test predicted choices within this space. They show that the majority of individuals select strategies that are boundedly optimal. Further, individual differences in what people choose to do are successfully predicted by the analysis. Two issues are discussed: (1) the performance of the minority of participants who did not find boundedly optimal adaptations, and (2) the possibility that individuals anticipate what, with practice, will become a bounded optimal strategy, rather, than what is boundedly optimal during training.

Keywords: Bounded optimality, adaptation, bounded rationality, constraints, utility maximisation.

It is known that people choose strategies that adaptively distribute memory and planning between internal and external resources according to the cost/benefit structure of the task environment (Payne, Howes & Reader, 2001; Gray, Sims, Fu, Schoelles, 2006; Marewski & Schooler, 2011). For example, it is known that lower action costs when solving the 8-puzzle decreased the amount of planning by participants (O'Hara and Payne, 1998). This can lead to longer solution paths and less learning in terms of transfer to other solution paths. Similarly, it is known that people make strategic use of computer help systems when the costs of accessing such systems are low but otherwise prefer strategies that rely on imperfect memory (Gray

and Fu, 2004). Gray et al. refer to people as preferring imperfect information in the head over perfect information in the world. Many others have demonstrated, or argued for, the adaptive nature of how people use the external task environment (Brumby, Howes & Salvucci, 2007; Cary & Carlson, 2001; Charman & Howes, 2003; Duggan & Payne, 2001; Edwards, 1965; Gigerenzer & Selten, 2001; Gigerenzer, Todd & the ABC Group, 1999; Gray, Sims, Fu & Schoelles, 2006; Kirsh & Maglio, 1994; Payne, Bettman & Johnson, 1993; Payne, Duggan & Neth, 2007; Payne, Richardson & Howes, 2001; Schönplflug, 1986; Smith, Lewis, Howes, Chu, & Green, 2008; Tseng & Howes, 2008; Walsh & Anderson, 2009).

The proposal that people distribute memory adaptively contrasts to the idea that people routinely offload cognitive processing (Hollan, Hutchins, & Kirsch, 2000). A weak version of the offloading hypothesis is that people simply make use of the environment to perform cognitive functions. The stronger version is that people favor the use of the environment over the use of internal psychological resources. Ballard, Hayhoe, Pook, and Rao (1997), for example, argued that people use a *minimal memory* strategy to copy arrangements of color blocks on a computer display. Participants in an experimental study tended to make frequent visual checks of the target pattern, rather than attempting to encode the pattern in memory. The idea that people favor offloading was rejected by S. J. Payne et al. (2001) in favor of a view of people as adaptive decision makers (J. W. Payne, Bettman & Johnson, 1993). The idea that people make use of the environment was not disputed, but rather Payne et al. (2001) questioned the idea that people *minimize* the use of internal cognitive resources, or that they are cognitively lazy. According to Payne et al. (2001), people choose to trade offloading with cognitive processing given the cost/benefit structure of the task. According to Gray et al. (2006), differences in temporal cost of just a few hundreds of milliseconds are enough to shift the allocation of resources from relying on the environment to more memory intensive strategies.

The purpose of the current article is to test whether an *individual's* selection of strategies for short-term remembering can be explained as *boundedly optimal* remembering. Behaviour is boundedly optimal if it can be predicted with a theory in which subjective utility is maximised given the bounds imposed by individual information processing capacities and their experience (Howes, Lewis & Vera, 2009; Howes, Lewis & Singh, 2014; Howes, Vera, Lewis & McCurdy, 2004; Lewis, Howes & Singh, 2014; Lewis, Vera & Howes, 2004). The hypothesis moves beyond previous

work on adaptation to consider the cost/benefit structure of the task environment in two respects. The first, is in the assumption that people do not merely adapt the distribution of memory but that they can also find boundedly optimal adaptations. The second, is in the assumption that the bounds are not only those of the task environment but that they are also due to an individual's own particular resource limits and experience.

The particular resource limits that we focus on in this article are memory limits. We are interested in the extent to which people are boundedly optimal given experience of their own performance on a simple short-term remembering task. Previous work, concerning what people choose to remember, has been conducted by Ballard, Hayhoe and Pelz (1995) and by Gray et al. (2006), amongst others. Following Ballard et al. (1995), Gray et al. (2006) used a Blocks World task to study the choices that people make about what to remember. The participant's task was to reproduce patterns of coloured squares (blocks) visible in a Target window, in a Workspace window. There might, for example, be 8 blocks, each of a different color, positioned randomly in a 4 x 4 grid. Gray et al. manipulated a lockout period, a period of 0 to 3000 milliseconds, before the target pattern became available after the participant moved the mouse over it. On average blocks encoded in memory per visit to the Target window increased from just over 2 to just under 3 blocks as the lockout period increased, demonstrating adaptation to external costs. The studies reported below use a variant of Gray et al.'s task to show that choice about how much to encode is not only adaptive but that it is also boundedly optimal.

Bounded Optimality

The motivations for this paper come from the bounded optimality framework proposed in Howes, Lewis and McCurdy (2004), Howes, Lewis and Vera (2009), Payne and Howes (2013) and Lewis, Howes and Singh (2014). The term "bounded optimality" was first used to refer to algorithms that maximize utility given a set of assumptions about problems and constraints in machine reasoning problems (Horvitz, 1988). According to Russell and Subramanian (1995), page 575: "an agent is boundedly optimal if its program is a solution to the constrained optimization problem presented by its architecture and the task environment." We assume here that individual, embodied, human minds correspond to the kinds of boundedly optimal machines defined by Russell and Subramanian (1995) although we do not, for the

moment, make the distinctions between the different types of bounded optimality articulated by these authors. Unlike, Russell and Subramanian (1995) our goal is not to develop the formal basis of Artificial Intelligence, but rather to test bounded optimality as a hypothesis about human behaviour. The key element that Russell and Subramanian (1995) bring is that rational behaviour is usefully defined as the deployment of optimal programs relative to constraints that include the cognitive architecture and experience (Lewis, Howes & Singh, 2014). In contrast, other approaches, more strongly influenced by economics, have tended to define rationality relative to the task environment (Anderson & Schooler, 1991) and/or in terms of sound principles of inference (Oaksford & Chater, 2007), though see Schooler and Anderson (1997) for a discussion of the relationship between rational analysis and processing bounds.

Bounded optimality is also influenced by key concepts in Reinforcement Learning (RL: Sutton & Barto, 1998; Dayan & Daw, 2008). Most importantly, Reinforcement Learning makes a commitment to the idea that learning what to do next concerns learning to maximise numerical reward signals through interaction with the environment. RL suggests extending Cognitive Science's traditional focus on goal-directed behaviour to a more explicit consideration of the utility of costs and benefits of interaction. Rather than merely describing goal states, RL demands that the value of states are considered with the aim of maximising the utility of behaviour to the agent. Our interest, therefore, is not with RL methods as hypotheses about the nature of human learning (e.g., Dayan & Daw, 2008), nor with RL methods as means of calculating optimal solutions (Chater, 2009), but rather with the formal definition of the RL problem as bounded utility maximisation. By implication, the problem of how to distribute memory is the problem of how to maximise reward signals through interaction with the environment.

In the following section, we report a number of examples of evidence for boundedly optimal behaviour. After that we contrast the optimization assumption required by bounded optimality with the explicit rejection of optimality found in Bounded Rationality (Gigerenzer & Selten, 2001; Simon, 1992).

Evidence for Bounded Optimality Given Response Variance

There is no empirical work to our knowledge that directly tests whether people are able to choose short-term memory strategies that are boundedly optimal.

However, there is relevant evidence in a range of perceptual-motor tasks. While these tasks do not demand that participants adapt what they choose to remember they do demand that people adapt movement strategies. A brief review is useful here because it will support a clearer articulation of the hypothesis. In particular, it will help us develop a theory of how the selection of remembering strategies is bounded by variation in how an individual performs a task, where variation is an inevitable consequence of internal constraints.

For example, Meyer, Abrams, Kornblum, Wright and Smith (1988), showed how a stochastic optimized-submovement model can explain simple movements. In the model, movements are described as an optimal compromise between the durations of primary and secondary submovements given noise on the control of movement caused by limitations of internal information processing and muscular control processes. The secondary movement acts to correct unintended, but inevitable, variance in the primary movement. Optimization is therefore bounded by internally generated variation in performance.

In an empirical investigation of Signal Detection Theory, Swets, Tanner and Birdsall (1961) tested the hypothesis that people select a boundedly optimal criterion. Participants were shown to select criterion levels, for the trade-off between correct detection and false alarms, that maximised utility. The optimization was achieved accounting for noise generated by the perceptual system in the signal level of targets and distractors. As with Meyer et al. (1988), people are boundedly optimal in the sense that they generate strategies that are optimal given bounds imposed by variation in human information processing mechanisms.

Trommershäuser, Maloney and Landy (2003) demonstrated that the assumption that participants were able to maximise financial gain in a task where they used a finger to point at a reward region and avoid a penalty region could be used to predict targeting. As with the previous examples, participants in this study learned to adjust where they pointed to their own particular profile of motor system noise. Participants who exhibited greater variation in the spread of where they pointed needed to adjust more in order to avoid the penalty region. Further studies have supported the idea that people learn boundedly optimal pointing strategies given variation in performance (Maloney & Mamassian, 2009; Trommershäuser, Maloney & Landy, 2003; Trommershäuser, Maloney & Landy, 2008).

Bounded optimality is also evident in more complex situations that require two responses and require those responses to be ordered (Howes, Lewis & Vera, 2009). By assuming that people were boundedly optimal in a series of Psychological Refractory Period (PRP) experiments, Howes et al. demonstrated that the interval between two responses could be precisely predicted for individual participants. Critically, the analysis defined optimality relative to the variance in the duration of each of the two responses. In order to maximise the utility of performance the model set the inter-response interval to a duration that was long enough to minimize the potential for response reversals without incurring a penalty for an excessive delay in the timing of the second response. The shorter the inter-response interval the greater the probability of a reversal error because of the variation in the duration of both responses. In other words, while participants cannot precisely predict the duration of one particular response, they can adjust performance to the response distributions. Howes et al.'s analysis of the PRP data showed that participants had made boundedly optimal adjustments to the duration between the two responses given individual characteristics of the response distributions.

Bounded optimality Versus bounded rationality

Bounded optimality shares much in common with bounded rationality. Bounded rationality is a framework for understanding behaviour that starts with the observation that people have limited time and limited capacities (Simon, 1997). These bounds impose limits on the extent to which people approximate the classical normative rationality that is, in contrast, insensitive to the reality of computation in the world. Bounded rationality also makes a commitment to the observation that behaviour often reflects adaptation to the structure of the environment (Gigerenzer, et al. 1999; Oaksford and Chater, 1994; Simon, 1997). In these respects, there is no difference between bounded optimality and bounded rationality.

Where Bounded optimality and bounded rationality differ is in the explicit rejection of optimality criteria that is evident in the definitions of bounded rationality provided by Simon (1997) and Gigerenzer and Selten (2001). While in earlier work, Simon pursued the idea that satisficing, a Bounded Rational heuristic method, was optimal for certain tasks (Kadane & Simon, 1977; Simon, 1955; Simon & Kadane, 1975), the predominant position articulated in his work was that the environment is too complex and computational resources are too limited for optimization to play a

role in explaining human behaviour (Simon, 1997). Gigerenzer, Todd and the ABC Research Group (1999) embraced this view of bounded rationality. Gigerenzer et al. work with the premise that much of human decision making and reasoning can be modelled with heuristics that “do not compute probabilities or utilities”. For Gigerenzer et al. the notion that people might optimize under constraints is a “demon”, a creature with unlimited capacity, that is rejected along with the “Unbounded Rationality demon” (pages 10-11; Gigerenzer et al., 1999). From Gigerenzer’s perspective optimization under constraints is paradoxical in that it seeks to explain limited information processing by assuming that the mind has essentially “unlimited time and knowledge”.

Bounded rationality and bounded optimality also differ sharply in practice. For example, Simon’s contributions to understanding short-term memory, long-term associative memory, and problem representations (e.g., see the compilations in Simon, 1979, 1989) were made without benefit of an explicit consideration of the effects of the utility functions that human participants might have adopted in the experimental situations. In contrast, bounded optimality requires consideration of utility functions (Howes, Lewis and Vera, 2009; Lewis, Howes and Singh, 2014).

However, our contention here is that in a wider range of tasks than previously thought, optimization algorithms can be usefully used to predict human behaviour. This is for three reasons. The first reason is the substantial recent literature, some of which is reviewed above showing that optimization can play a useful role in psychological theorising. The extensive repetition of more-or-less similar tasks, for example, gives opportunity, both in terms of time and knowledge, for optimal adaptation to occur on perceptual-motor tasks (see above) it may also do so on more complex, higher level decision making tasks, that involve constraints imposed by memory. The second reason is that these tasks, including many tasks used in experimental psychology, are what Savage (1954) called ‘small world’ tasks and it is therefore possible for the researcher to ascertain and solve the decision problem faced by the participants.

The third reason is that, in contrast to optimization under constraints, the cost of optimization is paid by the analyst, not by the participant. Bounded optimality does not assume that the mind is unlimited, rather it asserts that the analyst can usefully make use of optimization to test theories of the bounds (Lewis, Howes & Singh, 2014). This assumption is what Oaksford and Chater (1994) called ‘methodological

optimality'. A key benefit is that a prediction derived through optimization has a privileged status as an explanation for *why* people behave as they do because it allows a causal link to be established between bounds and behaviour (Hahn, 2014; Howes, Lewis & Vera, 2009; Payne & Howes, 2013).

Bounded optimality and probability matching

In contrast to the evidence, provided above, in favour of bounded optimality given response variance there are many studies that show that people probability match (Vulkan, 2002). Probability matching occurs when the frequency with which a choice is made is proportional to the probability that the choice maximizes subjective utility. Probability matching is often taken as evidence against the idea that people can be explained as performance optimisers. For a review of the probability matching literature see Vulkan (2000). While some studies have questioned the assertion that people probability match (e.g. Shanks et al., 2002), probability matching phenomena have been offered by others as evidence that people do not maximize subjective utility.

It is arguable whether people should probability match when they first experience a choice task. Indeed the boundedly optimal strategy for early stages of learning given choices with uncertain outcomes can be extremely difficult to ascertain. In general, the solution to these problems, depending on the assumptions, involves a period of exploration followed by convergence to the policy that exploits the highest rates of reward (Sutton & Barto, 1998; Gittins, 1989). In this paper we are interested in the bounded optimality of the strategies on which people converge after a period of exploration. In the studies, reported below we not only test whether participants are boundedly optimal but also whether they probability match. We ask which of these two theories is better able to explain the data.

Overview

If individuals are boundedly optimal then they should seek strategies that are optimal given subjective utility and bounds on short-term remembering. Each individual should not offload and should not make a minimal use of memory. They should not exhibit any bias in the use of memory away from what is measurably boundedly optimal for that individual. They should not continue to probability match

in cases where the prediction of maximized utility deviates from the prediction of probability matching.

In what follows, we report two experiments. In each experiment, participants are asked to make choices that have implications for the remembering strategies that can be deployed while performing a laboratory version of an everyday task. The choices concern how many items to hold in memory when copying messages from a calendar to an email system. Structurally, the task is similar to that deployed by Gray et al. (2006) which involved copying colour squares but, in contrast, to Gray et al. it uses a no-choice/choice paradigm so as to measure the utility of a range of different memory loads for each participant. Therefore, unlike for Gray et al. it is possible to draw conclusions about the efficiency of a participant's choice of memory load. The paradigm is described further in the next section.

Experiment 1

Experiment 1 was designed to test whether individuals used a boundedly optimal distribution of memory in a laboratory version of a memory task. The task involved copying appointments from a simulated 'email' application to a simulated 'calendar' application. Trials of the experiment were organised into a *no-choice* phase and a *choice* phase. This design is a novel variant on the choice/no-choice paradigm employed by Siegler and Lemaire (1997). In a no-choice/choice paradigm participants are first told which strategy to practice (a no-choice phase) and then asked to choose their preferred strategy (a choice phase). Siegler and Lemaire (1997) introduced the choice/no-choice paradigm, with the choice phase first, so as to address weaknesses with choice studies of adaptation. With the no-choice phase they were interested in obtaining unbiased estimates of the performance characteristics of a set of strategies, and in particular in recording the speed and accuracy of each strategy.

We reversed Siegler and Lemaire's (1997) choice/no-choice order so that the no-choice phase could act as a training phase ensuring that all participants were equally exposed to every strategy. This provided performance data that could be used to inform and evaluate strategy selection during the choice phase.

The purpose of the no-choice phase was to elicit a performance profile of a subset of the memory strategies available to participants. The space of strategies for the email-copying task encompasses variation along a number of dimensions, including number of items to encode, encoding method, and rehearsal method.

However, rather than explicitly elaborate a large space of strategies varying along these dimensions and instructing participants on the micro-structure of each strategy within this space, we presented participants with a sequence of trials that varied in the number of items that the participant was asked to remember. We used a small space of possible list lengths (3, 4, 5, 6, 7, 8, and 9) and instructed participants to attempt to remember the corresponding numbers of names. For example, a participant might be asked to remember 5 names and copy these to the 'calendar'. In addition, the incremental presentation of the list of appointments further restricted the encoding strategy. These list lengths and instructions, thereby, encouraged participants to adopt a strategy that involved the encoding of a certain number of appointments in memory. The participants' performance on each list length provides us with a measure of how utility varies along this single, but important, dimension of the space of strategies.

In order to test for bounded optimality it was important to provide an explicit and measurable utility regime for the participants. The goal for the participants was to copy a set number of items in as fast a time as possible. Utility for participants was therefore defined in terms of the time taken to copy all of the items. The faster that all items were copied, then the sooner the participants would be paid and could leave the laboratory. Importantly, we operationalized errors in terms of time. For example, in one condition, only correctly copied items counted towards the total number of items copied. Incorrect copies, for example recalling the wrong item, resulted in wasted time and a lower reward. As we describe later, participants were instructed to *correctly* copy n appointments as quickly as possible. Further, they were instructed that their choice of number of appointments to be presented/copied should be made to achieve this end. There was a trade-off between selecting strategies that appeared faster, in the absence of errors, and the increased risk of errors.

For the purpose of the analysis, as reported in the results, we defined utility in terms of the rate at which items were copied. 'Rate' refers to the number of items copied per second. We use rate as a measure of performance because it can vary as the participant progresses through the experimental trials. If participants are boundedly optimal then they should make remembering choices that maximize the rate at which items are copied.

Method

Participants

Forty native English speaking students from the University of Manchester participated in the study. They received £5 (\$8.09) as compensation for their time.

Materials

■ Following Gray et al. (2006), the task involves copying information from one computer application window to another. However, our task involved copying appointment information, where Gray et al.'s (2006) participants copied information about a spatial arrangement of colour blocks. A program was written in Microsoft Visual Basic 6 that simulated the email and calendar functions from Microsoft Outlook. This program ran on an ordinary personal computer with a keyboard and mouse.

To mimic the experience of receiving email, all visual elements of the original Outlook interface were reproduced. In addition, a single large button was included in the Inbox. The caption for this button was "Click for timeslots". Clicking on this button caused a message to be displayed in the box to the right of the button. This message was of the form "09:00: Appointment with NAME", where NAME was replaced with the name, in capitals, of the person at that appointment time. Each click of this button increased the time displayed by one hour and changed the name presented. Only one name and appointment was visible at a time (the display is illustrated schematically in the left panel of Figure 1).

Once all appointments had been displayed, a button in the bottom left hand corner of the screen labelled "Calendar" was enabled. Clicking this button changed the interface into a modified version of the calendar function from Outlook (illustrated schematically in the right panel of Figure 1). There were nine different boxes into which users could enter text. These boxes corresponded to the appointment times, thus the uppermost box was labelled 9.00, the second 10.00 and so on down to the bottom box labelled 17.00. Participants entered text into a box by clicking on it and typing using the keyboard. Pressing "Tab" cycled down through the boxes. Beneath these appointments there was a button labelled "Finished". All other buttons, menus and features of both the Email and Calendar interfaces were disabled. Every time the

participant clicked a button or entered text via the keyboard the program recorded and time-stamped the event.

A stimulus set of eight male and eight female first names (e.g. ROSE) was constructed. All of these to-be-remembered names were deemed familiar to native English speakers and were four letters long. Each name began with a different first letter.

Design and Procedure

Participants were divided into two groups of equal size. The cost of making an error (the payoff) was manipulated across the two groups; therefore, they were labelled “Low Error Cost” and “High Error Cost”. In the Low Error Cost condition each incorrectly copied appointment was counted as an error. In the High Error Cost condition all of the appointments in a trial were counted as errors if one or more of them was copied incorrectly. The experiment was divided into two phases; the No-choice phase followed by the Choice phase. Each phase was completed when participants had correctly copied a specified total number of appointments into the calendar.

All participants were instructed that they were required to copy appointments from the email application into the calendar. They were told that within each message there were two pieces of information: the name of the person to be met and the time of the appointment. However, they were also informed that the first appointment was always at 09.00 and all appointments were always one hour apart and in sequence, therefore, only the names and the order they were presented in needed to be remembered.

Appointments were presented in trials. On each trial, participants were required to view between 3 and 9 appointments before the calendar function was enabled and appointments could be copied across. The number of appointments that participants were required to read before copying across was an independent variable during the No-choice phase and a dependent variable during the Choice phase.

Blocks of seven trials were presented consecutively during the No-choice phase. Each trial within a block contained a different number of appointments to be copied. Therefore, each of the seven list lengths ranging from 3 appointments up to 9 appointments was represented once within each block. The order of trial presentation within each block was determined randomly. The order varied across blocks and

across participants. For every appointment, on every trial, the program randomly selected a name from the stimulus set of sixteen names. The only constraint on this process was that no name was allocated to more than one appointment on the same trial.

After completing a practice trial containing 3 appointments, all participants were asked to copy 200 appointments into the calendar as quickly as possible. It was emphasized that making errors was only problematic insofar as it slowed down the overall time taken and that their aim should be to finish as quickly as possible, rather than finish with as few errors as possible.

At the start of each trial a screen appeared indicating the total number of appointments remaining and the number of appointments that would be presented on that particular trial. When participants clicked a button labelled “OK” this screen was replaced with the email interface. Participants were presented with each of the appointment names in turn and then required to copy the names in uppercase letters into the appropriate slots within the calendar. After copying they were free to edit the text as much as desired and when satisfied should click the button labelled “Finished”. The program then provided feedback about the number of appointments correctly copied and highlighted in red any slots incorrectly completed. Any erroneous spellings, lowercase letters, or spaces left within a calendar slot when the “Finished” button was clicked, caused the item to be scored as an error. When the error feedback was provided another button was enabled that participants clicked to begin the next trial. Participants could not go back to correct errors, they could only progress to the next trial.

In the High Error Cost group, a single error on any of the appointments meant that all of the appointments from that trial were classed as errors and no points were awarded. Thus, if there were 8 appointments presented during a High Error Cost trial, and errors were made when copying 3 of them, then the overall total to-be-copied would have remained the same. In the Low Error Cost group all appointments correctly copied reduced the overall total to be copied. Thus, if there were 8 appointments presented during a trial and errors were made when copying 3 of them then the overall total to-be-copied would have been reduced by 5. In the analysis below, we refer to the reduction in the total number of items to-be-copied on a trial as the points achieved on the trial.

After 200 appointments had been correctly copied in the No-choice phase, participants received the instructions for the Choice phase. This phase was identical to the No-choice phase except that participants were allowed to select the number of appointments that were presented on each trial. This choice was implemented at the start of each trial by clicking on one of seven buttons labelled 3, 4, 5, 6, 7, 8 or 9, respectively. For the choice phase, participants were instructed to correctly copy a further 100 appointments as quickly as possible and that their choice of number of appointments to be presented/copied should be made to achieve this end.

The importance of the fact that participants had to *correctly* copy 100 appointments during the choice phase is worth restating. If a participant failed to correctly copy items then their target of items remaining to be copied was not reduced. As a consequence, unlike in many experiments, errors were not merely counted by the experimenter, but rather they had real consequences for the time taken by the participant.

Results

Average list length selected

Figure 2 is a plot of the mean rate at which items were copied against list length (number of items) for both the no-choice and the choice data. The mean rate at which items were copied was calculated according to the following procedure. For every trial, we recorded the trial duration, the selected length (3, 4, 5, 6, 7, 8 or 9 items), and the number of appointments correctly copied. The trial duration was defined as the interval between the end of the previous trial and the end of the current trial. This duration therefore included the time cost of moving from one trial to the next. For each participant and each list length (number of appointments), the rate R_k for a trial was then calculated by dividing the number of appointments copied by the amount of time taken for the trial. We then calculated an average rate for each participant and each list length.

In the low cost condition, a single point was awarded for each successfully copied item. For example, a participant who attempted to copy 5 items and made 1 error would get 4 points. In the high cost condition, a single point was awarded for each successfully copied item, unless there were any errors in which case no points

were awarded. For the same example, attempting to copy 5 items and making 1 error would result in no points. We defined an error as a failure to copy an item correctly.

As we have said, an important property of the rate is that error costs are reflected in the measure because when participants made errors it cost them time (by an amount contingent on the condition). We refer to the list length (number of appointments held in memory) associated with the highest rate of copies as the boundedly optimal list length. We assume that the boundedly optimal strategy involved the selection of the boundedly optimal list length.

Participants did not all complete the same number of trials because while they were required to copy 100 appointments during the choice phase, they were free to choose how many to copy on each trial.

In Figure 2 it can be seen that for the Low Error Cost condition the mean list length selected was 5.05 ($SD = .67$; Mode = 4.93, $SD = .69$) and for the High Error Cost condition the mean selection was 5.00 ($SD = .78$; Mode = 4.98, $SD = .80$). There was no statistically significant difference between the conditions for the means or modes ($ts < 1$). The absence of a difference in the choice phase list length is disappointing but, conversely, it can be seen in Figure 2 that the mean participant choice in both conditions is predicted by the no-choice phase rates.

Figure 2 gives the illusion that the rate for each list length was a point value when in fact they were distributions. This is made clear in Figure 3 which shows the frequency distribution of rate for each list length across all participants in both cost conditions and across both choice, and no-choice phases, of the experiment. Qualitatively, the figure suggests that some choice discriminations are relatively easy. It is easy to see that a list length of 9 is worse than a list length of 4. Other discriminations, e.g. between 4 and 5 are relatively difficult because of the overlap in the rate distributions.

Correlation between the boundedly optimal list length and the selected list length

In Figure 2 there appears to be a correspondence between the strategy with the highest rate in the no-choice phase and the chosen strategy in both conditions. In order to test this hypothesis further we first defined the mean boundedly optimal list length B_p for each participant p , as

$$B_p = \arg \max_{s \in S} R_{p,s}$$

Where s is one of the set of possible list lengths S and the rate of reward for a list length, $R_{p,s}$ is defined above. B_p is therefore defined as the list length s that maximised the rate R for participant p .

We pooled participants from both conditions and found a significant correlation between the boundedly optimal list length and the list length that participants actually selected, $r(38) = .35, p = 0.027$. Participants for whom it was predicted that they would take on larger list lengths did so, suggesting that the boundedly optimal list length predicted 12.25% of the variation between participants. This finding offers initial support for bounded optimality. Given the assumption that the boundedly optimal strategy involves the selection of the list length that allows each participant to maximise their own utility, then we know that the boundedly optimal strategy predicted by the theory is correlated with the list length actually selected by participants.

Probability matching

Before analysing the extent to which people were bounded optimal, we first wanted to reject the possibility that participants probability matched. The idea was that rather than using a strategy involving a list length that yielded the maximum utility, participants selected a list length in *proportion* to the probability that the strategy yielded the maximum utility, that is the highest rate of copies (e.g. see Shanks et al., 2002; Walsh & Anderson, 2009). We took the list length that each participant selected most frequently during the choice phase, called the *highest frequency* list length, and plotted the probability selected against the probability that it was the list length that maximised utility. If the participants were probability matching then we expected Figure 4 to show a straight line through 0,0 and 1,1. However, there was no significant correlation between the logit transformed probability selected and probability optimal ($r(38) = .061, p = .707$).

Frequency of Utility Maximisation

We tested whether the highest frequency list length selected by participants was selected more frequently than was predicted by probability matching. We first found the list length that was selected most frequently by each participant. We then found the probability that this list length was boundedly optimal for that individual. Recall that each list length has a distribution of rate (Figure 3) and so the probability that a list length is boundedly optimal is simply the probability that a sample of that list length's rate is greater than a sample of any other list length's rate. The data for both list length and probability were positively skewed and we therefore used a permutation test. A permutation test, with 10,000 resamples, contrasting probability selected and probability maximum utility, was significant $p < .001$. The mean probability boundedly optimal was 0.49 and the mean probability of selection of the most frequent list length was 0.80. Participants were significantly more likely to select their highest frequency list length than was predicted by probability matching (reflected in the fact that most of the data in Figure 4 are above the probability matching line).

Comparing boundedly optimal to suboptimal choice

We were interested in comparing the predictions of the boundedly optimal list length to list lengths that implied the encoding of fewer items in memory and to list lengths that involved encoding more items in memory. We examined the means of all list lengths with fewer items (optimal₋₁, optimal₋₂ etc.) and found that optimal₋₁ predicted as many selections, or more, than all others that had fewer items. The corresponding result was found for optimal₊₁. For this reason we focused these analyses on optimal₋₁ and optimal₊₁ (if optimal₋₁ performs worse than the boundedly optimal list length then optimal_{-n} will also perform worse). The optimal₋₁ list length offers a test of the offloading hypothesis; this is the hypothesis that people routinely offload to the environment. Contrasting the maximum utility list length (max) to optimal₋₁ and optimal₊₁ offers a test of the precision of the predictions. Figure 5 is a bar graph contrasting the average percentage of trials on which each of bounded optimal, optimal₋₁, and optimal₊₁ list lengths predicted participant performance. On

average, boundedly optimal predicted 55% ($SD = 32$) of participant selections, whereas optimal₋₁ and optimal₊₁ predicted 8% ($SD = 12$) and 17% ($SD = 24$) respectively.

A permutation test was used with 10,000 resamples to contrast the proportion of predicted selections in the choice phase for each of the three list lengths (bounded optimal, optimal₋₁, and optimal₊₁). The permutation test was used because the distributions for optimal₋₁ and optimal₊₁ were skewed. The boundedly optimal theory was a better predictor of selections than optimal₋₁ ($p < .001$), and a better predictor of selections than optimal₊₁ ($p < .001$). Optimal₋₁ and optimal₊₁ were equally poor predictors. All other strategies, for example, optimal₋₂, optimal₋₃, predicted even fewer selections.

Individual differences

We wanted to investigate individual differences across trials. For each participant and each trial, we computed the probability that a random use of any one list length would be better, that is deliver a higher rate of copies, than a random use of any of the other list lengths. The distribution of rates for each list length was set to the empirical distribution of rates for each list length for values over trials 1 to $k-1$. The computation of the probability was achieved using 1000 Monte Carlo trials for each list length on each trial of the experiment.

For example, consider a scenario in which there were only list lengths of 3 and 4. If participant 1 had experienced rates $R_3 = (0.4, 0.3, 0.4, 0.7)$ for list length 3 and rates of $R_4 = (0.2, 0.4, 0.6, 0.6, 0.5)$ for list length 4 then probabilities were calculated by sampling n pairs with replacement, with one element of each pair from R_3 and R_4 and then counting the frequency that the sample for 3 was greater than the sampled rate for 4. For example, if the sample generated from R_3 was 0.7 and from R_4 was 0.2 then the frequency that list length 3 was better than list length 4 would be incremented by 1. Once calculated, for each individual participant on each trial, this frequency was divided by the total number of sampled pairs, n , so as to generate the probability that each list length would generate a higher rate. In the analysis of the results presented below, rather than in this illustrative example, samples were taken from all 7 list lengths and the probabilities were calculated for each list length relative to all other list lengths.

The advantage of probability best, over rate (used previously), is that it is sensitive to the uncertainty in the rate associated with each strategy, as represented by the empirical distribution functions. Probability best is a measure of the likelihood that a strategy is boundedly optimal for the individual participant. One strategy, for example using a list length of 6-appointments, may be associated with a higher mean rate than a strategy using a 5-appointments list length but may also have much higher variation, or the two strategies may have such high variation that they are effectively indistinguishable; the probability best measure is sensitive to the distribution of rewards for each strategy.

We plotted the probability that utility was maximised with each list length. In Figure 6 each panel represents the likelihood that each list length maximised utility given a particular participant's trial-to-trial experience through the experiment.

Participants 8, 5, 4, 12, 13, and 14 were selected to represent the diversity of performance. In each panel, the no-choice phase is to the left of the vertical bar and the choice phase is to the right. Circles represent the selected list length. Each list length is represented with a different colour. We analysed all participants irrespective of condition.

Participants 4, 12, 13, and 14 (Figure 6) are presented because each selected the boundedly optimal list length on the majority of trials. In addition, each of these participants chose a different list length from the others and the figure, therefore, illustrates some of the individual differences in performance. For participant 4, a list length of 7 allowed them to maximise utility and the participant selected a list length of 7. For participant 12, list length 4 was boundedly optimal and the participant selected list length 4. For participant 13, list length 5 was boundedly optimal and the participant selected this list length on the majority of trials. For participant 14, list length 6 was the boundedly optimal and it was also selected. In addition, for participant 14, while list length 6 was not the boundedly optimal at the beginning of the choice phase, practice improved its performance to the extent that it became the boundedly optimal list length.

Participant 5 (Figure 6) was selected because there was no clear boundedly optimal list length. All strategies have probabilities below about 0.4 and three of the strategies (4, 5 and 6) have probabilities in a narrow range between 0.2 and 0.4. On some trials the participant selected list length 6 and on some list length 5, but these

strategies and list length 4 are indistinguishable (it is not clear that there is a distinct boundedly optimal list length).

Participant 8 (Figure 6) was selected because their behaviour illustrates choice phase performance that is not predicted by the theory. For this participant, by the end of the choice phase, the probability that list length 4 is the boundedly optimal list length is about 0.5 and the probability of all of the others is below 0.2. Despite the discrimination between the probabilities, the participant has selected a list length that is unlikely to allow them to maximise utility (list length 6) on the majority of choice phase trials.

Plots of the probability that each list length maximised utility for each participant are provided in Supplementary A.

Regression of selection against trial

We analysed whether participants were more likely to select the optimal list length with trial. We estimated the fixed effect of trial on whether, or not, bounded optimality predicted list length selection. A repeated measures logistic regression computed probability boundedly optimal for the selected list length against trial and revealed a significant positive slope ($p < 0.001$). Participants were more likely to select the boundedly optimal list length as trial progressed. Figure 7 displays a plot of the fit for each participant.

Discussion

The results offer support for the bounded optimality hypothesis.

- (1) As predicted, there was a positive correlation between the boundedly optimal list length and the selected list length; individuals who were predicted, on the basis of their measured performance across the strategy space, to choose a higher working memory load did, in fact, do so. While the magnitude of the errors points to variation, there is indication that strategy choice is sensitive to individual performance.
- (2) As predicted individuals were more likely to select the list length with the maximum utility than list lengths that involved encoding

more or fewer items in memory (see Figure 5). Further, a repeated measures logistic regression showed that participants were significantly more likely to select the boundedly optimal list length with practice.

Despite the positive evidence, a substantial portion of the data could not be accounted for as boundedly optimal choice of list length. For example, 8 participants became less likely to select the boundedly optimal strategy as trials progressed (they exhibit a negative slope in the regression reported in Figure 7). Four other participants persistently selected a list length that was not bounded optimal. (They exhibit a flat regression slope in Figure 7). We return to this result in the General Discussion. In addition, the manipulation of the external reward signal failed to generate a difference in either the predicted list length or in the list length selected by participants. These problems are addressed in the design of Experiment 2.

Experiment 2

While Experiment 1 offered some support for the hypothesis that individuals would choose to use bounded optimality list lengths, there was no effect of the manipulation of payoff function on the strategies selected by participants. Therefore, in Experiment 2, rather than manipulate the cost of an error, we manipulated the number of points awarded for a successful copy such that, in one condition, there was a greater incentive to copy larger list lengths (more details are given below).

Method

Participants

Twenty native English-speaking students from the University of Manchester participated in the study. They received £5 (\$8.09) as compensation for their time.

Design and Procedure

The goal for the participant was to score a set total of points by copying appointments into the appropriate slots in the calendar. As in the “Low Error Cost” condition of Experiment 1, a score for a trial was computed from the number of correctly copied appointments made when copying other appointments. Zero points were awarded for errors.

The key manipulation was the relationship between the number of appointments copied on a single trial and the number of points received for that trial. This was a between participant manipulation across two groups of equal size. In the “Linear” group participants received a single point for each appointment correctly copied. The total number of appointments to be copied in the Choice phase was doubled from Experiment 1, meaning participants had to score 200 points in both the No-choice and the Choice phases. In all other respects, the Linear condition was the same as the Low Error Cost condition from Experiment 1.

In the “Exponential” group, the number of points received for a trial increased exponentially according to the number of appointments correctly copied. Specifically, for copying 1 appointment participants received 1 point and the total trial points for each additional correctly copied appointment were 2, 3, 4, 7, 11, 17, 27 and 42. The target number of points in both the No-choice and the Choice phases was set at 310 points. This number was derived from the mean data from the Low Error Cost condition in Experiment 1. Assuming participants made the same number of errors on the same trials, then during the No-choice phase participants would take the same number of trials to reach 310 points in the Exponential condition as it took to reach 200 points in the Linear condition. This kept the amount of practice prior to the Choice phase approximately equivalent across both groups. Of course, these points totals did not necessarily result in both groups completing the same number of trials during the Choice phase – indeed the purpose of our manipulation is to produce a difference between the two groups.

At the start of the experiment all participants were given a table and graphic that outlined the relationship, between appointments copied and points scored, that was specific to their condition. It was emphasized to participants that they should aim to score the target points total as quickly as possible. All other aspects of the method were the same as in Experiment 1.

Results

Unless stated otherwise, all measures are computed and analysed in the same way as for Experiment 1, except that here *rate* refers to the number of points acquired per second rather than the number of items.

Later trials on which fewer than 10 participants contributed were excluded. The following analyses, therefore, use data from the no-choice phase and trials 1 to 39

of the choice phase. No other data was excluded from the analyses. The mean number of trials completed in the no-choice phase did not differ between the Exponential condition ($M = 41.10$, $SD = 16.40$) and the Linear condition [$M = 46.60$, $SD = 6.36$, $t(18) < 1$].

Average List length Selected

The mean list length that participants selected was larger in the Exponential condition ($M = 6.99$, $SD = 1.32$) than in the Linear condition [$M = 5.07$, $SD = .78$; $t(18) = 3.98$, $p = .001$, $d = 1.33$] supporting the hypothesis that people can adapt remembering strategies to the objective points-based utility function specified in the instructions. The average of each participant's mode list length produced the same significant difference [Exponential, $M = 7.20$, $SD = 1.62$; Linear, $M = 4.90$, $SD = .99$; $t(18) = 3.83$, $p = .002$, $d = 1.31$]. This relationship is illustrated in Figure 8, where the rate at which items were copied is plotted against the list length.

Correlation between the boundedly optimal list length and the selected list length

As with Experiment 1, we pooled participants from both conditions and found a significant correlation between the boundedly optimal list length and the list length that participants actually selected, $r(18) = .77$, $p < .001$. The RMSE was 1.34 and the boundedly optimal list length explained 59.29% of the variance. Participants for whom it was predicted that they would select larger list lengths did so, suggesting that boundedly optimal choice predicted a substantial part of the variation between participants. While the correlation does not tell us whether participants were biased, it does tell us that participants who were measurably able to copy larger list lengths did so.

Probability matching

As with Experiment 1, for each participant and each trial, we computed the probability that a random use of any one list length, and therefore strategy, would be better; that is would deliver a higher rate, than a random selection of any of the other list lengths. The computation was achieved using 1000 Monte Carlo trials for each list length on each trial of the experiment. Each list length was represented by the empirical distribution function formed from the values of its rate over trials 1 to $k-1$.

Figure 9 is a plot of each participant's most frequently selected list length. It provides a representation of the extent to which probability of selection was predicted by the probability that the selection was bounded optimal. If, on average, participants used probability matching then probability selected should match probability bounded optimal. The line of best fit should pass through 0,0 and 1,1. While there was a correlation [$r(18) = .62, p = .003$], a permutation test revealed that participants were significantly more likely to select the most frequent choice than predicted by probability matching [$p < .001$].

Individual differences

Inspection of the individual plots, of probability boundedly optimal versus trial, revealed a similar pattern of individual variation as that observed in Experiment 1. First, 12 of the 20 participants (5 in the Exponential condition, 7 in the Linear condition) selected the boundedly optimal list length on the majority of trials. 2 participants selected between a set of strategies all of which could have been the boundedly optimal, but which were essentially indistinguishable, and 6 participants systematically selected a list length that was not the predicted list length. Of this last group, 4 participants selected larger strategies than the bounded optimal, and 2 selected smaller strategies than predicted by bounded optimality (all within ± 2 of the bounded optimal). Plots of the likelihood that each list length maximised utility for each participant are provided in Supplementary B.

We visually inspected the response data file where all key-presses and mouse clicks were recorded. This log showed that for 49% of trials in the Exponential condition, during the recall phase, participants did not initially enter the complete names in each box. Instead, they selected each response box in turn and only entered the first letter of a name in each box. Once a letter had been entered in each box they then returned and entered the remaining letters of the name. This strategy was less frequently observed in the Linear condition (17% of trials), where participants entered the complete name in a box and rarely returned to a box subsequently. This strategy offered less benefit for the Linear condition as there was less reward for accurately remembering large list lengths. The mean list length selected was larger for the first letter strategy ($M = 6.65, SD = 1.70$) than the complete name strategy [$M = 5.46, SD = 1.01, t(23) = 2.10, p < .05, d = .80$].

Comparing boundedly optimal to suboptimal choice

We compared the maximum utility list length to a list length that involved one fewer items in memory (optimal₋₁) and a list length that involved encoding one more item (optimal₊₁). Figure 10 is a barplot of the percentage of trials on which each of the three strategies (bounded optimal, optimal₊₁ and optimal₋₁) predicted a participant's selection.

Permutation tests were used to contrast the proportion of predicted selections in the choice phase for each list length. (The distributions for optimal₋₁ and optimal₊₁ were positively skewed.) The boundedly optimal strategy was a better predictor of selections than optimal₋₁ ($p = .003$), and a better predictor of selections than optimal₊₁ ($p = .004$). Neither optimal₋₁, nor optimal₊₁, was a better predictor than the other. On average boundedly optimal predicted 46% ($SD = 33$) of participant selections, whereas, optimal₋₁ and optimal₊₁ predicted 12% ($SD = 13$) and 14% ($SD = 17$), respectively. We used a repeated measures logistic regression to test whether each theory – bounded optimal, optimal₋₁, and optimal₊₁ – predicted more, or fewer, participant selections with trial. We found no effect of trial on whether optimal predicted the choice $p = 0.1285$. We did find an effect of trial on whether optimal₋₁ predicted the choice $p = 0.006$. There was also an effect of trial on whether optimal₊₁ predicted the choice $p = 0.0391$. Both optimal₋₁ and optimal₊₁ become significantly worse at predicting the participant's choice.

Discussion

In Experiment 2, half of the participants received exponentially increasing rewards for those list lengths, and therefore those strategies, that required more memory. The other half received linearly increasing reward. As predicted, individuals in the exponential condition selected significantly larger list lengths than individuals who received linearly increasing rewards, demonstrating that participants can adapt choice of memory strategy to utility. Evidence that participants not only adapted, but were also boundedly optimal was also present. The boundedly optimal list length was a significantly better predictor than either optimal₋₁ or optimal₊₁ supporting the idea that participants were boundedly optimal. However, it was also the case that many participants failed to select the predicted list length.

General Discussion

Two experiments used the no-choice/choice paradigm to test the hypothesis that individuals can make boundedly optimal choices when remembering items for short periods of time. The no-choice phase of the experimental paradigm allowed us to empirically measure performance on a range of strategies and, thereby, calculate the boundedly optimal strategy for each individual. The choice phase allowed us to test the prediction that people would not only adapt but that they would do so by choosing a list length, and therefore a strategy, that maximized utility. The findings (Experiment 2) are consistent with previous findings (Gray et al., 2006) that people are able to adapt their use of memory; on average people choose to remember a different number of items depending on the payoff regime. In addition, both Experiments 1 and 2 offered evidence that adaptations of the majority of participants were bounded optimal. In both experiments, the boundedly optimal strategy offered significantly better predictions of average performance than strategies with fewer items, or more items, than the boundedly optimal strategy – suggesting that the hypothesis that people minimize the use of memory (Ballard, Hayhoe, Pook, & Rao, 1997; Hollan, Hutchins, & Kirsch, 2000) is inconsistent with the evidence and further supporting the hypothesis that people are adaptive to costs and benefits (Payne, et al., 2001; Gray et al., 2006). Further, in Experiment 1 regression analysis indicated that with practice participants became significantly more likely to select the boundedly optimal strategy as they experienced more trials. Correlations between optimal and selected for each individual suggest that in both experiments the majority of participants adapted to their own individual performance characteristics. The individual differences between these participants were therefore not merely described but *predicted* by the bounded optimality analysis.

The validity of these findings is contingent on the effectiveness of the no-choice/choice utility learning paradigm (Siegler & Lemaire, 1997; Walsh & Anderson, 2009) which allowed us to determine the utility of strategies other than that chosen by the participants. Validity was also contingent on the fact that participants were asked to maximise an explicit utility function. Errors were operationalized in terms of time. To the extent that the results showed which participants were bounded optimal, they did so given a paradigm in which utility, and therefore optimality,

involved a quantifiable speed/accuracy trade-off. People can, it appears, adjust what they choose to remember over short time periods so as to maximise utility given speed/accuracy constraints, at least, they did so in the reported studies.

The Value for the Current Work

Experiments 1 and 2 go beyond previous work (e.g. Payne et al. 2001; Gray et al., 2006) in three important respects. First, the experiments add support for the idea that the majority of people can, not only, adapt their use of memory but in addition they can adapt to just the right extent. On the whole, if a participant could achieve their highest rate with a list length of say 5 then this is the list length that they used when given a choice. No previous experiments requiring people to remember items for short time periods has demonstrated that behaviour is substantially consistent with a theory that demands boundedly optimal adaptation. Second, the results show that these participants maximized the rate at which items were copied by choosing an *individually appropriate* list length. The correlations between boundedly optimal list length and chosen list length in both experiments show that participants who copied items at a higher rate with a particular list length chose that list length during the choice phase of the experiment.

Third, the results show that some participants failed to choose a boundedly optimal list length. The fact that the experience of some of these individuals led to no clear boundedly optimal list length suggests one explanation, but other participants failed to choose what the analysis shows was a clear bounded optimum. We discuss, below, the implications of this apparent form of ‘suboptimality’ and its relationship to the findings of Fu and Gray (2004).

Future Work

Explaining behaviour that was not bounded optimal

There were 13 (out of 60) participants in the two studies who persistently selected a list length, and therefore a strategy, that was not bounded optimal, for example participant 8 in Figure 6. Visual inspection of the probability boundedly optimal for each list length, as presented in Figure 6, suggests that given the evidence available to these participants, they should not have been unsure about which was best, yet they persistently failed to select this list length. If there was a clear

boundedly optimal strategy then exploration of suboptimal list lengths should have been unnecessary. These participants did not all select a larger, or all select a smaller list length; although 10 of the 13 participants selected a list length that was larger, usually by 1 more memory item, than that associated with the boundedly optimal strategy.

One explanation for the behaviour of these 13 participants is that they were somehow less able than others to determine the relative utility of different list lengths; in other words, it is plausible that they simply failed to appreciate the correct utility ranking. Just as some participants were less able to remember items, so some may have been less able to determine the relative utility of remembering more or fewer items. If this is the case then it is possible that these participants are boundedly optimal given their utility discrimination capacity. However, further studies are required to test this hypothesis.

Another possible explanation is that participants believed that practising sub-optimal strategies would make the strategies optimal. Many of the participants who did eventually achieve a boundedly optimal remembering strategy, did so by practising a list length that was initially sub-optimal. Practice both improved the performance of the strategy and reduced uncertainty about its performance. For example, see participant 4 in Figure 6. Again, further evidence is required.

Lastly, it is possible that people exhibit stable suboptimalities (Fu & Gray, 2004). Evidence reported by Fu and Gray (2004) who studied users of computer applications suggests that the preferred, less efficient procedures, have two characteristics: (i) the preferred procedure is well practised and can be deployed for a variety of task environments, and (ii) the preferred procedure has a structure that gives step-by-step feedback on progress, or in other words, it is more interactive. According to Fu and Gray (2004) these participants are suboptimal because they are biased to use more interactive and general procedures. This bias towards procedures that are globally efficient leads people to exhibit stable local suboptimalities. However, Payne and Howes (2013) point out that “any conclusion of suboptimality is relative to a particular theory of utility, and local suboptimalities may well be globally optimal. The challenge is to find a theory of utility, context (global or local), and mechanism that explains the observed behaviour. One aspect of such an approach would involve a systematic exploration the implications of different theories of subjective reward (Singh, Lewis, Barto & Sorg, 2010; Janssen & Gray, 2010). Ultimately suboptimal

adaptation to memory must be explained. The character of the explanation, we anticipate, will have the form: "people were not adapting to X but to Y." (p. 76)

Explaining exploration

When people learn a new task, over repeated trials, they engage in both exploratory and exploitative behaviours. They must sometimes choose strategies in order to exploit knowledge about likely rewards and they must sometimes choose strategies in order to explore what the rewards are for each strategy (Cohen, McClure & Yu, 2007; Sutton & Barto, 1998). Indeed, exploration is one benefit of probability matching. For the most part, the studies reported in the current article focused on how people exploit the knowledge that they have gained during a no-choice phase, which might be described as a forced exploration of the strategy space. More specifically, the focus was on how, during the choice phase, people exploit the knowledge that they have gained on previous trials.

While our analysis focused on exploitation, it is evident that participants may have engaged in some exploratory behaviour, at least at the beginning of the choice phase. Regression analysis of the Experiment 1 data showed that participants were significantly less likely to select the boundedly optimal strategy at the beginning of the choice phase than toward the end. Further, analyses suggested that probability matching did not do well at explaining how exploration/exploitation was managed (Figures 3 and 8). A fuller analysis of the observed exploratory behaviour might test an optimal data selection theory of which strategies people choose to explore (Oaksford & Chater, 1994, 2003; Oaksford & Wakefield, 2003; Nelson, 2005, 2008; Lelis & Howes, 2011). For example, it might be the case that on earlier choice trials, when the performance of each strategy is still relatively unclear, that participants choose a strategy so as to maximise gain in information, rather than to maximise immediate reward. One possibility is that participants in our experiments operationalized the value of information in terms of the extent that it facilitated discrimination between the alternative memory strategies. Another possibility is that they operationalized value as the expected gain in choice utility obtained by a likely choice reversal (assuming that when not deliberately exploring they would exploit the boundedly optimal choice). See Lelis and Howes (2011) for a discussion.

Discriminating between these theories in the utility learning paradigm that we have investigated above is beyond the scope of the current article, but the no-choice/choice paradigm may be useful in the future. The key strength of the paradigm – that it exposes the distribution of the reward for each strategy in the strategy space – should allow a-priori prediction of the information gain from each choice.

Acknowledgements

This work was partly supported by the Office of Naval Research and by a development grant from the Faculty of Humanities, University of Manchester. The authors are particularly grateful to Niels Taatgen, Julian Marewski, Leendert van Maanen, Cvetomir Dimov, Chris Janssen, Sashank Varma and Wai-Tat Fu for their invaluable comments on an earlier manuscript.

References

- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological science*, 2(6), 396-408.
- Ballard, D. H., Hayhoe, M. M., & Pelz, J. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, 7, (1), 66-80.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioural & Brain Sciences*, 20, 723-767.
- Brumby, D.P., Howes, A., & Salvucci, D.D. (2007). A cognitive constraint model of dual-task trade-offs in a highly dynamic driving task. In S.J. Payne and B. Begole (Eds.) *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2007)* (pp. 233-242). New York, NY: ACM Press.
- Cary, M., & Carlson, R. A. (2001). Distributed working memory resources during problem solving. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 27, 3, 836-848.
- Charman, S. C., & Howes, A. (2003). The adaptive user: an investigation into the cognitive and task constraints on the generation of new methods. *Journal of Experimental Psychology. Applied*, 9(4), 236-48.
- Chater, N. (2009). Rational and mechanistic perspectives on reinforcement learning. *Cognition*, 113(3), 350-64.

- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 1481, 933-942.
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, 31(3), 187-276.
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective & Behavioural Neuroscience*, 8, 429-453.
- Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience & Biobehavioral Reviews*, 30(2), 239-271.
- Duggan, G. B., & Payne, S. J. (2001). Interleaving reading and acting while following procedural instructions. *Journal of Experimental Psychology: Applied*, 7, 297—307.
- Edwards, D. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*, 2, 312-329.
- Gigerenzer, G., & Selten, R. (2001). Rethinking rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 1–13). Cambridge, MA: MIT Press.
- Gigerenzer, G., Todd, P. M. & the ABC Research Group. (1999). *Simple heuristics that make us smart*. NY: Oxford University Press.
- Gittins, J. (1989). *Multi-armed bandit allocation indices*. NY: Wiley.
- Gray, W. D., & Fu, W. (2004). Soft constraints in interactive behaviour: the case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28, 359-382.
- Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behaviour. *Psychological Review*, 113, 461–482.
- Hahn, U. (2014). Experiential limitation in judgment and decision. *Topics in Cognitive Science*, 6(2), 229-244.

- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7, 174 - 196.
- Horvitz, E. J. (1988). Reasoning about beliefs and actions under computational resource constraints. In T. Levitt, J. Lemmmer, & L. Kanal, (Eds.), *Uncertainty in Artificial Intelligence 3*. North Holland: Amsterdam.
- Howes, A., Vera, A., Lewis, R. L. & McCurdy, M. (2004). Cognitive constraint modeling: A formal approach to reasoning about behaviour. In K. D. Forbus, D. Gentner & T. Regier (Eds.), *Proceedings of the 26th Annual Meeting of the Cognitive Science Society*.(pp. 595-600). Hillsdale, NJ: Erlbaum.
- 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, 4, 717-751.
- Howes, A., Lewis, R. L., & Singh, S. (2014). Utility maximization and bounds on human information processing. *Topics in Cognitive Science*, 6(2), 198-203.
- Janssen, C. P., & Gray, W. D. (2012). When, What, and How Much to Reward in Reinforcement Learning- Based Models of Cognition. *Cognitive science*, 36(2), 333-358.
- Kadane, J. B., & Simon, H. A. (1977). Optimal strategies for a class of constrained sequential problems. *The Annals of Statistics*, 5, 237-255.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513- 549.
- Lelis, S. & Howes, A. (2011). Informing decisions: How people use online rating information to make choices. *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems*, 2285-2294.
- Lewis, R. L., Howes, A., & Singh, S. (2014). Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in Cognitive Science*, 6(2), 279-311.
- Lewis, R.L. & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29, 3, 375-419.
- Lewis, R. L., Vera, A., & Howes, A. (2004). A constraint-based approach to understanding the composition of skill. In M. Lovett, C. Schunn, C. Lebiere, & P. Munro (Eds.) *Proceedings of the Sixth International Conference on Cognitive Modeling*. (pp. 148-153). Mahwah, NJ: Lawrence Erlbaum.

- Maloney, L. T., & Mamassian, P. (2009). Bayesian decision theory as a model of human visual perception: Testing Bayesian transfer. *Visual Neuroscience*, 26, 147-155.
- Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: an ecological model of strategy selection. *Psychological Review*, 118(3), 393-437.
- Meyer, D. E., Abrams, R. A., Kornblum, S., Wirght, C. E., & Smith, J. E. K. (1988). Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review*, 95, 340-370.
- Nelson, J. (2005). Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, 112 979-999.
- Nelson, J. (2008). Towards a rational theory of human information acquisition. In M. Oaksford & N. Chater (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 143-163). Oxford: Oxford University Press.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 608-631.
- Oaksford, M., & Chater, N. (2003). Optimal data selection: Revision, review, and reevaluation. *Psychonomic Bulletin & Review*, 10, 289-318.
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford: Oxford University Press.
- Oaksford, M., & Wakefield, M. (2003). Data selection and natural sampling - Probabilities do matter. *Memory & Cognition*, 31, 143-154.
- O'Hara, K. & Payne, S. J. (1998). The effects of operator cost of planfulness of problem solving and learning. *Cognitive Psychology*, 35, 34-70.
- Payne, J.W., Bettman, J.R. & Johnson, E.J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- 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 (3), 370-388.
- Payne, S. J., Howes, A., & Reader, W. R. (2001). Adaptively distributing cognition: A decision-making perspective on human-computer interaction. *Behaviour and Information Technology*, 20, 339-346.
- Payne, S. J., Richardson, J., & Howes, A. (2000). Strategic use of familiarity in display-based problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1685-1701.

- Payne, S. J., & Howes, A. (2013). Adaptive interaction: A utility maximization approach to understanding human interaction with technology. *Synthesis Lectures on Human-Centered Informatics*, 6(1), 1-111.
- Russell, S. J., & Subramanian, D. (1995). Provably bounded optimal agents. *Journal of Artificial Intelligence Research*, 2, 575-609.
- Savage, J. (1954). *The Foundations of Statistics*. New York: Wiley.
- Schönpflug, W. (1986). The trade-off between internal and external information storage. *Journal of Memory and Language*, 25, 657-675.
- Schooler, L. J., & Anderson, J. R. (1997). The role of process in the rational analysis of memory. *Cognitive Psychology*, 32(3), 219-250.
- Shanks, D.R., Tunney, R.J., & McCarthy, J.D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioural Decision Making*, 15, 233-250.
- Siegler, R.S. & Lemaire, P. (1997). Older and younger adults' strategy choices in multiplication: testing predictions of ASCM using the choice/no-choice method. *Journal of Experimental Psychology: General*, 126, 1, 71-92.
- Simon, H.A. (1955). A behavioural model of rational choice. *The Quarterly Journal of Economics*, 69, 1, 99-118.
- Simon, H. A. (1992). What is an “explanation” of behaviour? *Psychological Science*, 3, 150-161.
- Simon, H. A. (1997). *Models of bounded rationality: Vol. 3. Empirically Grounded Economic Reason*. Cambridge, MA: MIT Press.
- Simon, H. A., & Kadane, J. B. (1975). Optimal problem solving search: All or none solutions. *Artificial Intelligence*, 6, 235-247.
- Smith, M. R., Lewis, R. L., Howes, A., Chu, A., & Green, C. (2008). More than 8,192 ways to skin a cat: Modeling behaviour in multidimensional strategy spaces. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1441-1446). Austin, TX: Cognitive Science Society.
- Singh, S., Lewis, R. L., Barto, A. G., & Sorg, J. (2010). Intrinsically motivated reinforcement learning: An evolutionary perspective. *Autonomous Mental Development, IEEE Transactions on*, 2(2), 70-82.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

- Swets, J. A., Tanner, W. P., & Birdsall, T. G. (1961). Decision processes in perception. *Psychological Review*, 68, 301-340.
- Trommershäuser, J., Maloney, L. T. & Landy, M. S. (2003). Statistical decision theory and tradeoffs in motor response. *Spatial Vision*, 16, 255-275.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Science*, 12(8), 291-297.
- Tseng, Y.-C., & Howes, A. (2008). The adaptation of visual search strategy to expected information gain. *Proceedings of the Twenty-Sixth Annual Conference on Human Factors in Computing Systems* (pp. 1075-1084). New York, NY: ACM Press.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14(1), 101-118.
- Walsh, M. M., & Anderson, J. R. (2009). The strategic nature of changing your mind. *Cognitive Psychology*, 58(3), 416-40.

Figure 1. Experiment 1. The experimental apparatus. On each trial participants were first presented with the 'email display'. They clicked the 'click for time slots' button until all appointments had been shown. The 'calendar' button then became available and pressing it caused the display to change to 'calendar display'. They then entered the names that they could remember into the time slots and pressed 'finished'.

Figure 2. Experiment 1: Mean rate at which items were copied for each list length in the no-choice phase and for the average list length chosen in the choice phase. Error bars are the 95% confidence interval for the mean chosen list length.

Figure 3. Experiment 1: The frequency distributions of the rate at which items could be copied with each list length. Data is for all participants in both conditions ($n=40$) and for both no-choice and choice phases.

Figure 4. Experiment 1: Probability selected versus probability bounded optimal for each participant's most frequently used list length. Probability matching predicts a straight line regression through 0,0 and 1,1 – which is not supported by these data.

Figure 5. Experiment 1: Percentage of predicted choice phase selections for the bounded optimal, optimal₋₁, optimal₊₁, and the selected strategy. Error bars are the 95% confidence interval for each strategy.

Figure 6. Experiment 1. Six panels that give illustrative examples of individual performance across trials. The probability that a strategy was the bounded optimal strategy is plotted against trial (see the text for a description of how this probability was calculated). Each strategy is represented by a different color. The selected strategy is represented by a circle. Participant 8 (top left) failed to find the bounded optimal strategy. Participant 5 (top right) did not exhibit a distinct bounded optimal strategy. Participant 4 (middle left) initially practiced a strategy lower than the optimal (strategy 6) before persistently selecting the bounded optimal strategy (strategy 7). Participant 12 (middle right) persistently selected the bounded optimal strategy (strategy 4) but also explored a higher memory strategy (strategy 5). Participant 13 (bottom left) persistently selected the bounded optimal strategy (strategy 5). Participant 14 (bottom right) practiced a strategy that became the bounded optimal.

Figure 7. Experiment 1. Plots of repeated measures logistic regressions of probability optimal selection (y-axis) against trial (x-axis) for each individual participant. Each plot indicates the probability that a participant selected the optimal list length with trial. (No axis labels are provided because of the number of plots.)

Figure 8. Experiment 2: Mean rate at which items were copied for each list length in the no-choice phase and for the average list length chosen in the choice phase. Error bars are the 95% confidence interval for the mean chosen list length.

Figure 9. Experiment 2. Probability selected versus probability bounded optimal for each participants most frequently used strategy. Probability matching predicts a straight line regression through 0,0 and 1,1. While there is a significant correlation [$r(18) = .62, p = .003$], probability bounded optimal and probability selected are significantly different [$V = 20, p < .001$].

Figure 10. Experiment 2: Percentage of choice phase selections predicted by the bounded optimal, optimal_{+1} , optimal_{-1} and the selected strategy against trial (choice phase only). Error bars are the 95% confidence interval for each strategy.

Author Manuscript

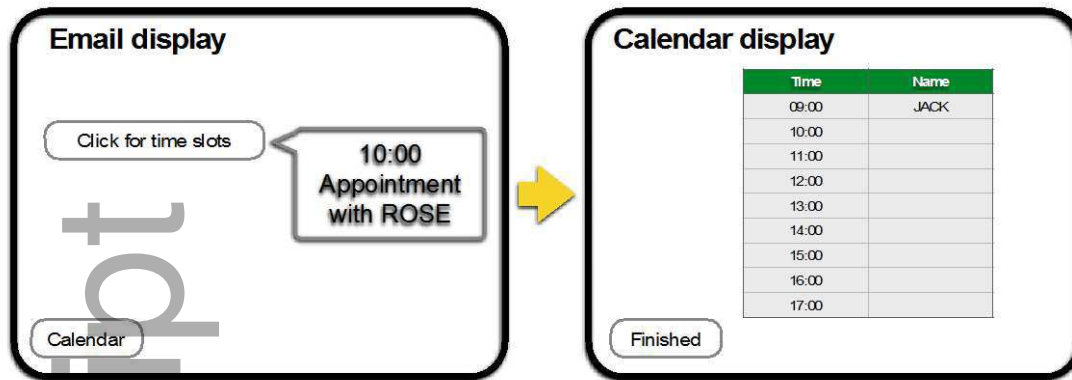


Figure 2

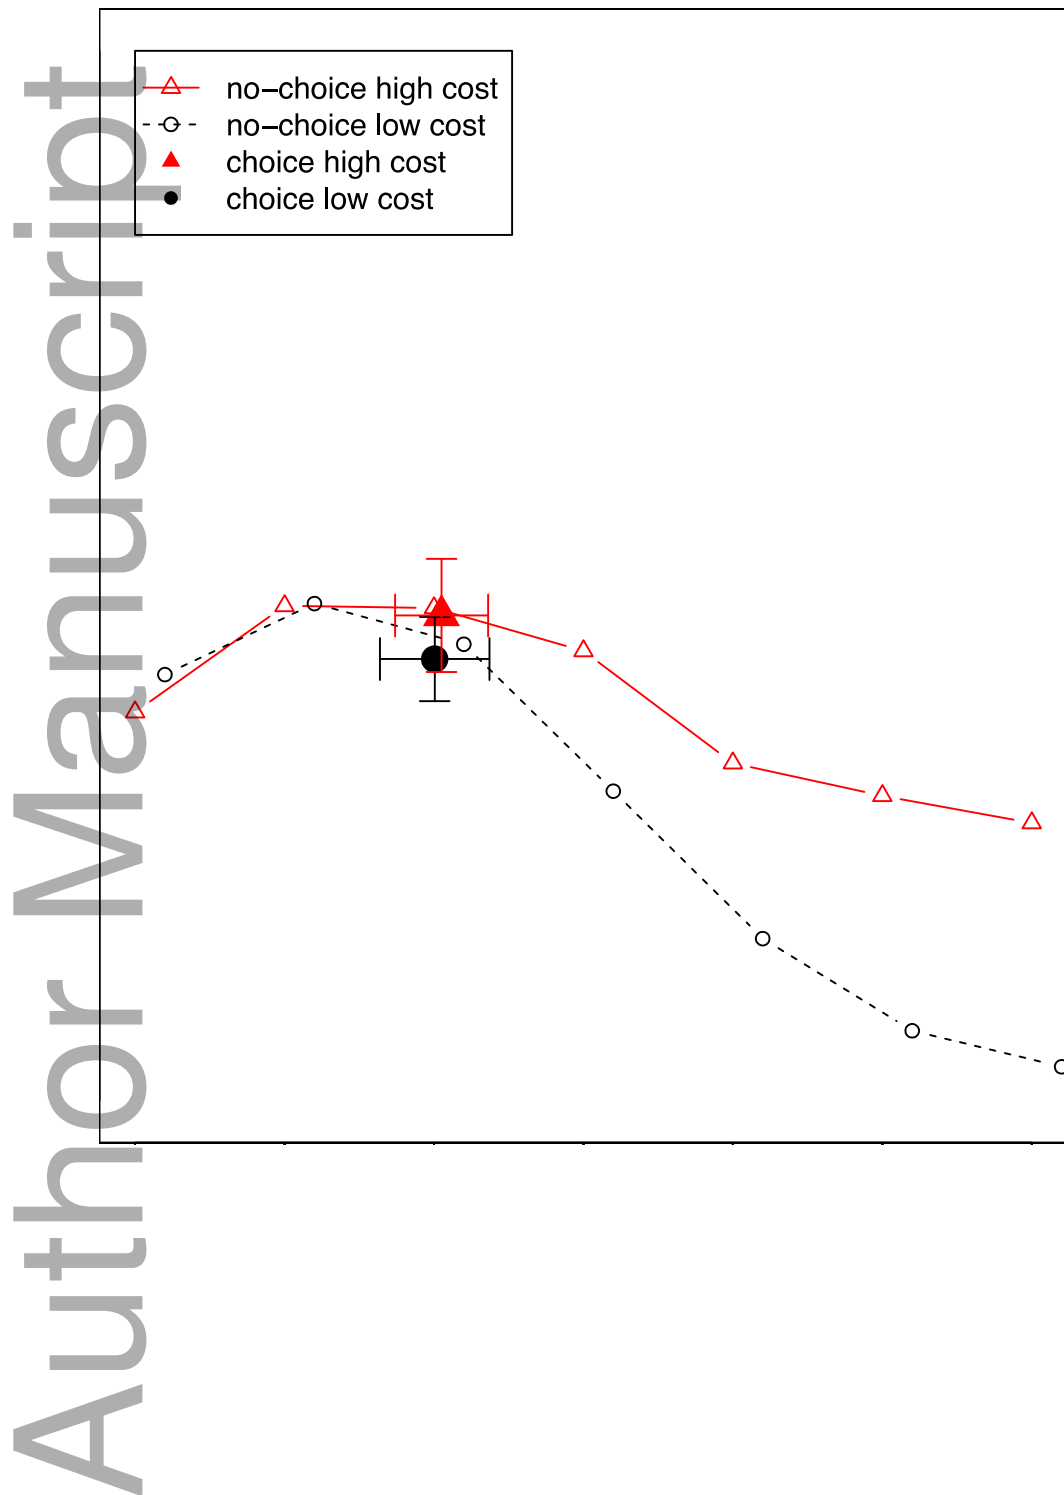


Figure 3



Figure 4




Figure 5




Figure 6

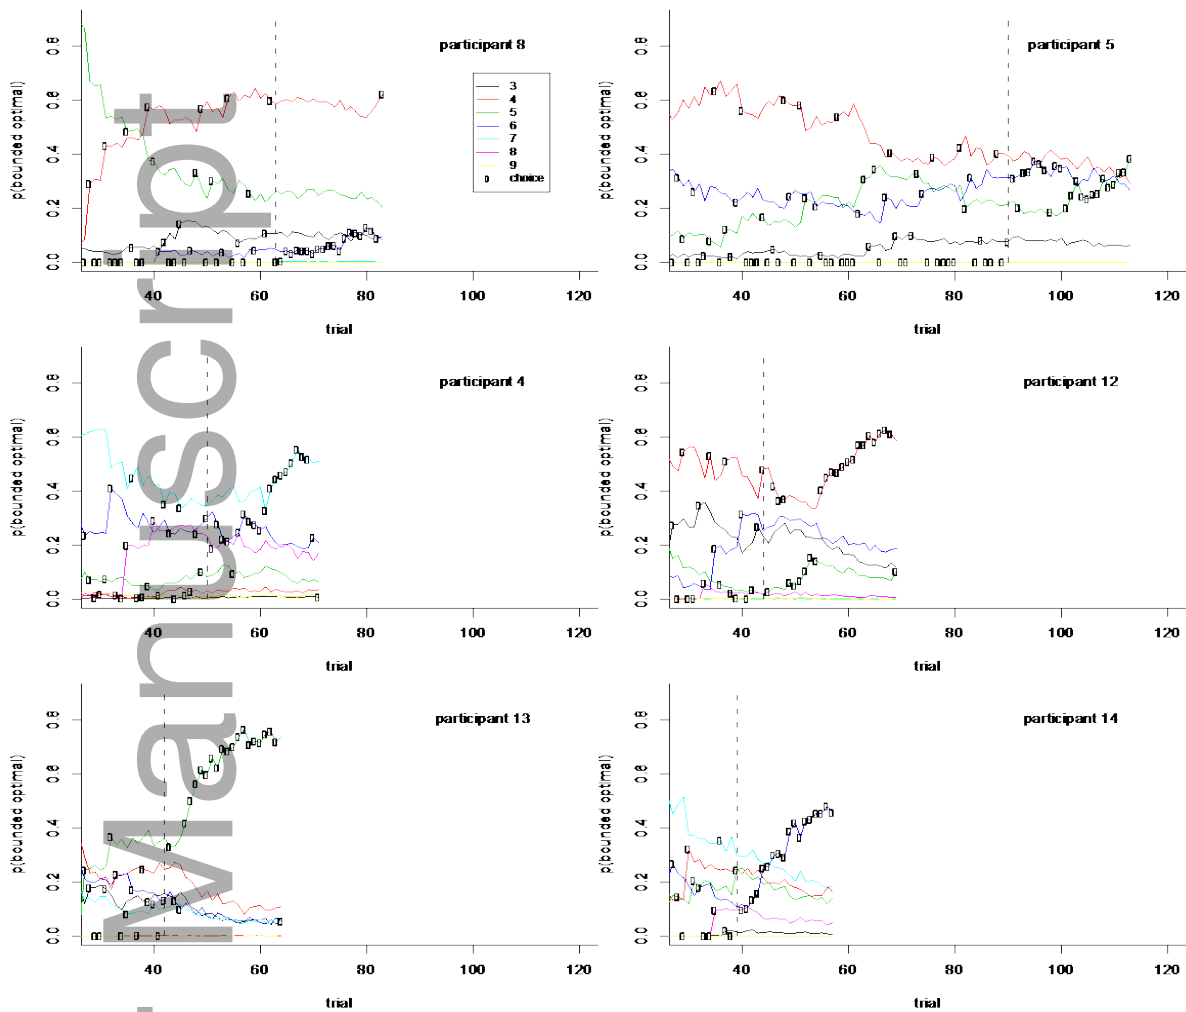



Figure 7.



cogs12271-fig-0007.docx

probability bounded optimal

Figure 8




Figure 9




Figure 10

