

The Benefits of Epistemic Action Outweigh the Costs

Paul P. Maglio (pmaglio@almaden.ibm.com)

IBM Almaden Research Center
San Jose, California

Michael J. Wenger (mwenger1@nd.edu)

Department of Psychology
University of Notre Dame

Angelina M. Copeland (ahill1@nd.edu)

Department of Psychology
University of Notre Dame

Abstract

Epistemic actions are physical actions people take to simplify internal problem solving rather than to move closer to an external goal. When playing the video game Tetris, for instance, experts routinely rotate falling shapes more than is strictly needed to place the shapes. Maglio and Kirsh (Kirsh & Maglio, 1994; Maglio, 1995) proposed that such actions might serve the purpose of priming memory by external means, reducing the need for internal computation (e.g., by way of mental rotation). This proposal requires that information provided by epistemic actions (e.g., additional views of the shape) serve the same function as memory primes, and that the benefit of such priming exceed the costs of performing the epistemic action. To calculate benefit, we used a novel statistical method for mapping reaction-time data onto an estimate of the increase in individual processing capacity afforded by seeing shapes in multiple orientations. To calculate cost, we used an empirical estimate of time needed to take action in a Tetris game. We found that the benefits of extra previews far outweigh the costs of taking extra action.

Introduction

Intuitively, expertise is associated with economy, efficiency, and optimality of performance (Logan, 1988; Newell & Rosenbloom, 1981). Data gathered from players of the video game Tetris, however, suggest that experts in fact take more action and do more backtracking in that game than do beginners (Kirsh & Maglio, 1994; Maglio & Kirsh, 1996). In Tetris, players maneuver two-dimensional shapes as they descend from the top of the screen to fit the shapes onto a playing surface at the bottom of the screen (see Figure 1). As players develop expertise, their performance becomes faster, yet they also *increase* the number of rotations made, requiring backtracking to correct for over-rotation.

Maglio and Kirsh (Kirsh & Maglio, 1994; Maglio, 1995) suggested that such counter-intuitive behavior might result from expert players taking actions not for their effect on the external state of the game, but for their effect on the player's internal problem-solving state. Such actions are called *epistemic actions*, and are used to simplify internal computations rather than to move closer to an external goal state. In the case

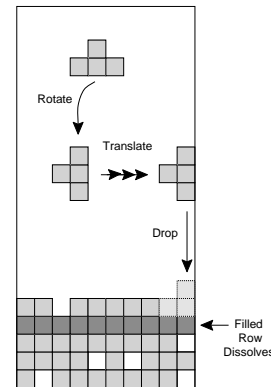


Figure 1: In Tetris, two-dimensional shapes fall one at a time from the top of the screen, landing on the bottom or on top of shapes that have already landed. The object of the game is to fill rows of squares all the way across. Filled rows dissolve and all unfilled rows above move down.

of Tetris, Maglio and Kirsh suggested that extra rotations provide information that would otherwise be obtained by mental rotation. Because shape identification can be facilitated by primes or previews oriented differently from a target shape (Cooper, Schacter, Ballesteros & Moore, 1992; Lassaline & Logan, 1993; Srinivas, 1995), it is plausible that Tetris players might rotate shapes to facilitate identification. Moreover, studies have shown that problem solving in other tasks is facilitated by taking extra actions, including counting coins (Kirsh, 1995), arranging letters to spell words (Maglio, Matlock, Raphaely, Chernicky & Kirsh, 1999), and solving the Tower of Hanoi (Neth & Payne, 2002).

Though both appealing and compelling, Maglio and Kirsh's **Epistemic Action Hypothesis** has not yet been put to a systematic test. The present study is the last in a series aimed at making such a test. To do this, we incorporate a novel method for evaluating costs and benefits of epistemic action. A strong test of the hypothesis that extra rotations in Tetris act as epistemic actions must account for two main predictions:

Priming. Previews of a Tetris piece can prime later decisions about that piece (e.g., whether it will fit into the board in its present orientation).

Net Benefit. Previews allow the player to do more with the available time than no previews (e.g., to consider additional alternative placements).

With respect to the **Priming** prediction, we previously reported experiments investigating the conditions under which previews aid performance: specifically, by reducing the response time (RT) needed to decide whether a Tetris piece fits into a given configuration of the playing surface (Maglio & Wenger, 2000, 2002). The results show that (a) increasing the number of distinct views that are available as previews leads to decreases in RT, and (b) the positive effect of the previews diminishes as the time between preview and decision decreases. The first of these findings is consistent with the possible shape-identification function of priming. The second can be understood by noting that, as the time for the decision draws closer, perceptual information available in the game display may be sufficient for the decision. Alternatively, with very short lags between preview and decision, it may not be possible for the prime to produce a memory retrieval that can be completed before a decision must be made.

The present paper investigates the **Net Benefit** prediction. Does the benefit obtained by priming with multiple orientations outweigh the costs of taking extra actions? That is, do extra previews actually allow the player to be more effective? If it were possible to measure a player’s ability to process game information per unit of play (imagine an instantaneous measure of a player’s “bandwidth” or capacity for internal computation), that measure would be higher when previews are available than when they are not. This would be true because the preview would relieve the player of the need to perform the rotation mentally, allowing more information (e.g., about the current state of the game) to be processed in a unit of time. But because epistemic actions come with a cost—that associated with making additional rotations and consequent corrections—the benefit of the epistemic action in increased capacity must be greater than the cost in increased playing time. Otherwise, experts would take longer to perform the task than would novices.

Simply put, to measure the cost of epistemic action, we will estimate the time cost of making extra rotations. This can be done by estimating the time required for the epistemic action and its correction (two keystrokes). To measure the benefit of epistemic action, we will estimate the increase in the player’s mental capacity given extra previews of the piece. For capacity, we need a measure that estimates the extent to which a system is capable of completing a given amount of work (computation) in a unit of time. To provide this measure, we draw on theoretical and empirical results of Townsend and colleagues (Townsend & Ashby, 1978, 1983; Townsend & Nozawa,

1995, Wenger & Townsend, 2000), which suggest that a specific method of characterizing the RT *distribution* provides a performance measure that can be interpreted as capacity. That measure is known as the *hazard function* of the RT distribution.

The Hazard Function

The hazard function is a *conditional probability function* that assesses the instantaneous likelihood of completing a process, conditional on not yet having completed the process. In this sense, it is distinct from the unconditional probability of completing a process, since this latter measure may be low for some given value of time, whereas the former may be very high. For example, although the unconditional probability of taking a long time to complete a process may be low, the probability of completing the process at that point in time given that the process has not yet finished may be high. The conditional character of this function allows it to be interpreted in terms of *intensity*, and indeed the function is known in certain engineering applications as the intensity function (Townsend & Ashby, 1978). That is, a processor that is operating at a high level of intensity will have a higher conditional probability of completing the processing task in the next instant (given that it has not completed its task yet) than will a processor that is operating at a low level of intensity. Mathematically, the hazard function is defined as

$$\begin{aligned} h(t) &= \lim_{\Delta t \rightarrow 0} \frac{P(t \leq T \leq t + \Delta t | T \geq t)}{\Delta t} \\ &= \frac{f(t)}{S(t)}, \end{aligned}$$

where $f(t)$ is the probability density function (in our case, for task completion times), and $S(t) = 1 - F(t)$ is the survivor function. Because the survivor function is the complement of the cumulative distribution function $F(t) = P(T \leq t)$, it gives the probability that the task has not been completed by time t , that is $P(T > t)$.

We can represent the hypothesis of an increase in capacity in terms of the relationship between the hazard function for conditions in which previews (available from the additional rotation) are present to those in which previews are absent. Specifically, if we let $h_p(t)$ be the hazard function when previews are present, and $h_n(t)$ be the hazard function when no previews are present, we can test the statistical hypothesis of

$$\frac{h_p(t)}{h_n(t)} > 1$$

This is easily and directly done using a set of well-understood log-linear regression methods (Collett, 1994; Cox, 1972). Although these methods have not traditionally been used with RT data, recent work suggests that they are robust with RT data (Wenger & Gibson, 2002; Wenger, Schuster, Petersen & Petersen,

2002 in press; Wenger, Schuster & Townsend, 2002). These procedures produce estimates of the magnitude of the proportionality relationships among a set of hazard functions, as well as tests of the hypotheses that these magnitude estimates are reliably different from 1. Finally, a simple transformation of these estimates produces a value that can be interpreted in terms of the percentage of change associated with a manipulation (e.g., the presence of a preview).

Costs and Benefits of Epistemic Actions

The hazard-function estimates let us discover how much the processing capacity is increased given a preview (i.e., the benefit of seeing additional views of the piece). But a complete account also requires an estimate of the cost of obtaining a preview (i.e., how long it takes to perform an extra rotation and then another rotation to correct for this extra move). To estimate the time required to do this, we required participants in the present study to make an additional keystroke at various unexpected points during task performance. We then multiplied this value by two to estimate the time required for both an additional rotation and its subsequent correction. The cost of these actions can then be expressed as a percentage of change in the total time for an experimental trial. Thus, we have both cost and benefit measured on the same scale. To the extent that epistemic actions work as hypothesized, the benefit (in percentage increase in capacity) should exceed the cost (in percentage increase in trial duration).

Method



To test this hypothesis, participants viewed a highly-restricted version of Tetris—the same as was used in our previous work (Maglio & Wenger, 2000, 2002). In this version, individual Tetris pieces are presented descending from the top of the computer screen toward the bottom. We will refer to the presentation of each piece in this sequence as a frame. At the end of the sequence, a test piece was presented immediately above a board configuration, and participants had to judge whether the piece in the test frame would or would not fit into the board, either in its current orientation or with a single rotation. In some of the trials of the sequence that preceded the test frame, preview information (about the test piece, the board, or both) was present. These trials were used to assess the potential benefit in capacity associated with the presence of a prime. In addition, we included a set of trials in which the test frame was replaced with the instruction to press one of two keys. These trials were randomly inserted into the total sequence of experimental trials, and data from these trials were used to estimate the cost of epistemic actions.

Participants

Fifteen participants were recruited from the University of Notre Dame's psychology subject pool, and participated voluntarily in exchange for course credit. All

participants had normal or corrected-to-normal vision, and none had participated in any of our previous Tetris experiments.

Design

Two kinds of experimental trials were used. The first was constructed so as to mimic the conditions that revealed differential effects of preview information in our earlier work (Maglio & Wenger, 2000, 2002). We will refer to these trials as the *Tetris trials*. The experimental design for these was a 2 (preview: present, absent) \times 2 (piece type: , ) \times 3 (relation between preview and test: prime related to board, prime related to test piece, prime related to board and piece) \times 2 (prime orientation with respect to test: same, different) \times 3 (lag between presentation of preview and test: 0, 1, 2 frames) \times 2 (status of the test piece with respect to the board: fit, does not fit) complete factorial design, with all factors manipulated within observers. The second kind of trials (*keypress trials*) involved a random assignment of these factors to a total of 100 trials that involved a sequence of pieces terminating in an instruction for a specific keypress, rather than presentation of a test piece and board.

Materials

All pieces and boards were constructed from 20 \times 20 pixel squares. Squares were outlined by light gray lines, 1 pixel in width, and were filled in solid black. The background for all displays was also solid black. All piece types were composed of four blocks. All boards were six blocks in height and width. Four "fit" boards were defined for each piece type, corresponding to four ways in which the piece could be snugly placed. Each such board was used with equal frequency. Materials were displayed on a 43 cm (diagonal) VGA monitor controlled by a PC-compatible computer. Onset and offset of each display was synchronized to the monitor's vertical scan. A standard keyboard was used to collect and time (to ± 1 ms) responses.

Procedure

Participants were run in a single session that lasted approximately 90 min. All sessions were conducted in a darkened room, with participants seated at an unconstrained distance from the monitor, and began with a five min period for dark adaptation. Participants were told that, on each trial, they would see a sequence of Tetris pieces presented very rapidly. The pieces in the sequence would begin falling from a location near the top of the screen. Each successive piece would appear below the one before to create a sequence of falling pieces, much as in the Tetris video game. Each piece was present for 250 ms, and each sequence consisted of between five and seven pieces, with the actual number determined randomly (and with equal likelihood) on each trial. In the Tetris trials, this sequence was terminated by the presentation of a test piece immediately above a board, and participants were instructed to indicate whether the test piece would fit

in the board (either without a rotation or with one rotation). “Fit” responses were given using the index finger of the observer’s dominant hand, and “no fit” responses were given using the index finger of the observer’s non-dominant hand, using the Z and / keys. In the keypress trials, the sequence of falling pieces was terminated with an instruction to press either the Z or / keys, with the specific key being chosen randomly on each instance of a keypress trial. In both the Tetris and keypress trials, participants were instructed to emphasize both speed and accuracy.

Results

We begin by examining the data from the Tetris trials to determine whether preview (or prime) information speeded responding, and if so, under what conditions the benefit was most pronounced. We do this by examining both mean RTs and hazard functions. We then consider the extent to which any benefits observed at the level of the hazard functions exceeded the costs associated with the keypresses. A criterion level of 0.05 was used in all analyses.

Mean RT

Our first question was whether the presence of a preview speeded responding relative to the absence of a preview. Mean RT for the preview-present trials was 528 ms, and the mean RT for the preview-absent trials was 673 ms, with this difference being reliable, $t = 5.09, p < .0001$. Our initial examination of the data from the preview-present trials indicated no main effects or interactions associated with piece type. Consequently, we analyzed these data using a 3 (relation between preview and test: prime related to board, prime related to test piece, prime related to board and piece) \times 2 (prime orientation with respect to test piece or board: same, different) \times 3 (lag between presentation of preview and test: 0, 1, 2 frames) \times 2 (status of test piece with respect to the board: fit, does not fit) repeated measures analysis of variance (ANOVA). The reliable results from this analysis are presented in Table 1. Mean RTs for the interaction involving prime orientation, fit status, and lag are shown in Figure 2.

The results here replicate those obtained in our previous work (Maglio & Wenger, 2000, 2002). When the preview primed both the test piece and the board, it produced faster RTs than when it primed either alone. RTs on trials where the test piece fit the board were faster than RTs on trials where the test piece did not fit. Finally, the positive impact of a preview increased with increases in lag between preview and test.

Overall Capacity Analyses

Results of the RT analyses showed positive effects of previews on performance, consistent with our previous work. As noted, inferences from these analyses must be limited to overall speed, and do not address processing capacity. To assess the extent to which previews produced increases in processing capacity, we fit log-linear regression models for the hazard functions to

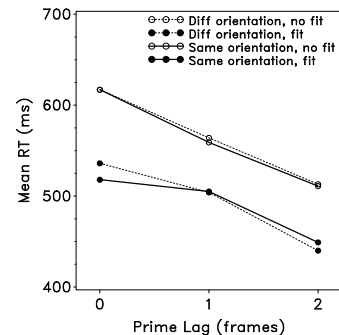


Figure 2: Mean RTs from the Tetris trials as a function of the relationship between prime and test frame (same or different), fit status of the test piece, and the lag between prime and test frame.

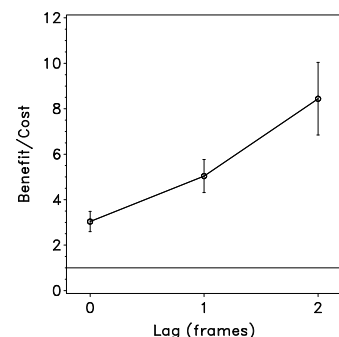


Figure 3: Mean value of the ratio of benefit to cost as a function of lag. Error bars represent 95% confidence intervals on the mean.

the data, allowing for the same set of main effects and interactions as allowed by the ANOVA. This analysis indicated that only a subset of the total possible effects were reliable. Table 2 presents the results for the model including only reliable effects. Having a preview increased capacity by almost 40% relative to not having a preview. This effect was attenuated at shorter lags, increasing as lag increased. When a prime was present, its effect was best when it was in the same orientation as the test piece, the board, or both.

Cost-Benefit Analysis

The overall analysis of the hazard functions showed that the previews do more than simply increase the speed of responding: they serve to increase the instantaneous capacity of players. To show that the benefit obtained from previews exceeds the cost of performing action, we must weigh the capacity increase against the time cost of taking extra actions. The keystroke trials allow us to estimate this cost: For each participant, we calculated mean RT on keystroke trials. We then doubled this value, and added the result to the fixed time associated with each trial (4.4 s). The cost estimate was this total divided by the fixed time, and

Table 1: Reliable results from the ANOVA on correct RTs from the Tetris trials.

Effect	<i>df</i>	<i>F</i>	<i>MSE</i>	<i>p</i>
Prime/test frame (PTF)	2	65.39	1192.90	< .0001
Fit status (FS)	1	91.05	7580.00	< .0001
Prime/test lag (PTL)	2	284.47	1404.89	< .0001
PTF × FS	2	38.93	1089.08	< .0001
FS × L	2	14.31	910.84	< .0001
Prime orientation × FS × L	2	4.65	591.01	0.0181

Table 2: Results of the regression analyses of the hazard functions on RT, Tetris trials.

Effect	<i>df</i>	β	χ^2	<i>p</i>	% Change
Prime presence/absence (PA)	1	0.3204	96.85	< .0001	37.8
Lag × PA	1	0.2896	295.40	< .0001	33.6
Prime orientation × PA	1	0.1087	80.20	< .0001	11.5

multiplied by 100 to obtain a percentage.¹

The benefit of previews was estimated as follows. For each participant, we fit the proportional hazards model, including only reliable effects, to obtain the percent change (as in Table 2). To test the hypothesis that benefit exceeded cost, we first subtracted the cost estimate for each subject from the benefit in each condition, and then subjected this difference score to a repeated measures ANOVA. At each lag, the ratio of benefit to cost was reliably greater than 1, and the value of this ratio increased with increases in lag, $F_{(2,28)} = 31.25$, $MSE = 10.75$ (see Figure 3). We can conclude not only that the benefit outweighs the cost of an epistemic action, but also that the benefit is greatest when memory load is greatest, namely with greatest time between preview and decision.

Discussion

A strong test of the hypothesis that epistemic actions can explain the extra rotations performed by expert Tetris players required that we determine whether the benefit of previews outweigh the cost of taking the action needed to create the previews. We approached this question of the **Net Benefit** in two steps. First, we predicted that presentation of previews should lead to an increase in capacity. Second, we predicted that increase in capacity should outweigh the time cost of performing the actions. To test the first prediction, we used the hazard function on the RT distribution as a measure of capacity. This conditional probability function can be interpreted in terms of the intensity with which a system can perform its processing tasks, and increases can be interpreted in terms of increases in capacity, something that is not possible with mean RTs (Townsend, 1990). Our results showed that previews produced increases in capacity, with benefits increasing as lag between preview and decision increased.

¹This cost overestimates the time of Tetris experts, who have likely internalized the epistemic strategy and therefore produce faster responses in practice.

To test the second prediction, we needed to estimate the cost of making an epistemic action. We made this estimate by having subjects perform simple keystrokes at unexpected intervals. We then doubled the time required for these actions to estimate the total cost of an epistemic action—an over-rotation and a correction. We then compared the increase in capacity (for each participant) to the cost, and found that the benefits more than outweighed the costs in all cases, with the ratio of benefit to cost increasing with increasing lag. Thus, this work provides strong evidence for the plausibility of Maglio and Kirsh’s **Epistemic Action Hypothesis** (Kirsh & Maglio, 1994; Maglio, 1995).

Recently, others have considered related models and shown similar results, largely in the context of human-computer interaction. Wright, Fields and Harrison (2000) describe the resources model, in which epistemic actions taken in the task environment are viewed as resources that an agent can rely on in accomplishing goals. Gray and Fu (2000) analyzed the task of setting a videorecorder in several conditions, and found that people routinely make tradeoffs in time and accuracy depending on configuration of the task environment. Neth and Payne (2002)’s study of the Tower of Hanoi problem showed epistemic actions benefit steady-state behavior but are helpful during learning. Kirlik (1998) develops a specific mathematical formulation of the epistemic action hypothesis, and tests in the context of expert short-order cooks. In all these models and studies, the costs of taking physical action in the external task environment is weighed against the benefit of the information gained by taking the action. Of course, not all actions are worth taking.

Our previous reports (Maglio & Wenger, 2000, 2002) explored conditions under which previews in Tetris improve performance. Our present results are consistent with these studies, and extend the findings by demonstrating that the benefits of taking action to generate extra previews outweigh the behavioral costs. Taken together, we see these results as being systematic,

thorough evidence in favor of the hypothesis proposed by Maglio and Kirsh. We also see the sum of these studies as providing a rich set of targets for a computational model, one that can allow for exploration of mechanisms and conditions under which epistemic actions might aid in skilled performance more generally.

References

- Collett, D. (1994). *Modeling survival data in medical research*. London: Chapman and Hall.
- Cooper, L. A., Schacter, D. L., Ballesteros, S., & Moore, C. (1992). Priming and recognition of transformed three-dimensional objects: Effects of size and reflection. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *18*, 43–57.
- Cox, D. R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society, B34*, 187–220.
- Gray, W. D. & Fu, W. (2000). The influence of source and cost of information access on correct and errorful interactive behavior. In *Twenty-second Annual Conference of the Cognitive Science Society*, pages 663–668, Mahwah, NJ: Erlbaum.
- Kirlik, A. (1998). The ecological expert: Acting to create information to guide action. In *Fourth Symposium on Human Interaction with Complex Systems*, pages 15–27. IEEE Computer Society.
- Kirsh, D. (1995). Complementary strategies: Why we use our hands when we think. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ: Erlbaum.
- Kirsh, D. & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, *18*, 513–549.
- Lassaline, M. E. & Logan, G. D. (1993). Memory-based automaticity in the discrimination of visual numerosity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 561–581.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.
- Maglio, P. P. (1995). *The computational basis of interactive skill*. Doctoral dissertation, University of California, San Diego.
- Maglio, P. P. & Kirsh, D. (1996). Epistemic action increases with skill. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*, pages 391–396, Mahwah, NJ: LEA.
- Maglio, P. P., Matlock, T., Raphaely, D., Chernicky, B., & Kirsh, D. (1999). Interactive skill in Scrabble. In *Proceedings of the Twenty-first Annual Conference of the Cognitive Science Society*, pages 326–330, Mahwah, NJ: Erlbaum.
- Maglio, P. P. & Wenger, M. J. (2000). Two views are better than one: Epistemic actions may prime. In *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, Mahwah, NJ: Erlbaum.
- Maglio, P. P. & Wenger, M. J. (2002). On the potential of epistemic actions for self-cuing: Multiple orientations can prime 2D shape recognition and use. In *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*, pages 608–613, Mahwah, NJ: Erlbaum.
- Neth, H. & Payne, S. J. (2002). Thinking by doing: Epistemic Actions in the Tower of Hanoi. In *Proceedings of the Twenty-fourth Annual Conference of the Cognitive Science Society*, pages 691–696, Mahwah, NJ: Erlbaum.
- Newell, A. & Rosenbloom, P. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*, pages 1–55. Hillsdale, NJ: Erlbaum.
- Srinivas, K. (1995). Representation of rotated objects in explicit and implicit memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *21*, 1019–1036.
- Townsend, J. T. (1990). Truth and consequences of ordinal differences in statistical distributions: Toward a theory of hierarchical inference. *Psychological Bulletin*, *108*, 551–567.
- Townsend, J. T. & Ashby, F. G. (1978). Methods of modeling capacity in simple processing systems. In J. Castellan & F. Restle (Eds.), *Cognitive theory (vol. 3)*, pages 200–239. Hillsdale, NJ: Erlbaum.
- Townsend, J. T. & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. Cambridge: Cambridge University press.
- Townsend, J. T. & Nozawa, G. (1995). On the spatio-temporal properties of elementary perception: An investigation of parallel, serial, and coactive theories. *Journal of Mathematical Psychology*, *39*, 321–359.
- Wenger, M. J. & Gibson, B. S. (2002). Using hazard functions to assess changes in processing capacity in an attentional cuing paradigm. under review.
- Wenger, M. J., Schuster, C., Petersen, L. M., & Petersen, R. C. (2002, in press). Application of proportional hazards models to reaction time data. In C. Bergeman & S. Boker (Eds.), *Notre Dame Series on Quantitative Methodologies*, pages xx–xx. Mahwah, NJ: Erlbaum.
- Wenger, M. J., Schuster, C., & Townsend, J. T. (2002). Testing capacity hypotheses: Applying proportional hazards and frailty models. Paper presented at the 35th Annual Meeting of the Society for Mathematical Psychology, Miami University of Ohio, Oxford, Ohio.
- Wenger, M. J. & Townsend, J. T. (2000). Basic response time tools for studying general processing capacity in attention, perception, and cognition. *Journal of General Psychology*, *127*, 67–99.
- Wright, P., Fields, R., & Harrison, M. (2000). Analysing human-computer interaction as Distributed Cognition: The resources model. *Human Computer Interaction*, *51*, 1–41.