



# Cognitive constraints on decision making under uncertainty

Christian Lebiere\* and John R. Anderson

Department of Psychology, Carnegie Mellon University, Pittsburgh, PA, USA

\*Correspondence: cl@cmu.edu

## INTRODUCTION

Models of decision making under uncertainty should be grounded in general cognitive processes reflecting pervasive constraints from the nature of our environment. Developing integrated models applicable across different tasks provides converging constraints that increase the predictiveness of models to new situations.

Decision making is such a rich discipline that it is often considered in relative isolation, leading to entire fields devoted to specialized aspects and domains. Decision making under uncertainty can be better understood through the prism of general theories of cognition, constrained by representations and mechanisms developed to account for the much broader range of human activities (Anderson and Lebiere, 2003). This argument is an elaboration of Simon's bounded rationality (e.g., Simon, 1957) approach to constraining the rationality of optimal decision making by the cognitive limitations of the decision maker. Further, those cognitive limitations, and more generally the entire decision making process, should be modeled in a computational framework that captures in detail not only the cognitive mechanisms and representations involved (Newell, 1990) but also motivational processes (Kruglanski et al., 2007; Kruglanski and Gigerenzer, 2011) as well as perceptual (especially attentional) and motor processes (Card et al., 1983) to reflect the constraints of the task environment. Finally, decision making involves not simply raw cognitive processes but also knowledge and strategies on how to approach the problem (Gigerenzer et al., 1999). Fortunately, unified theories of cognition enable the representation of declarative and procedural knowledge constructs in a way that interacts with the constraints of the cognitive and perceptual processes to provide a rich account of performance in the task. Providing a detailed but unified computational account of those factors and their interaction across a wide range of tasks is essential for a deeper understanding of

human decision making under uncertainty, as it involves general cognitive processes that are not limited to specific paradigms but take place across all human activities.

We illustrate those points by briefly describing a number of instances of our recent line of research in models of decision making. In particular, we want to highlight the importance of applying the same modeling approach to widely different paradigms of decision making (including domains that are not usually considered part of decision making) in order to bring the maximum force of converging constraints onto the problem. Indeed, the main issue with many decision making tasks is not that they are too difficult to model, but instead that too many distinct models provide roughly equivalent accounts of the data, making it hard to determine which provide a fundamental understanding of human decision making processes and which are merely well-fitted parametric descriptions of human performance.

## MODELS

Our initial model of decision making was applied to a task paradigm that is not traditionally considered part of decision making. Sequence-learning (e.g., Curran and Keele, 1993) usually involves speeded reaction tasks intended to investigate the impact of implicit learning processes on the detection of event sequences. In those tasks, a participant is exposed to a sequence of stimuli appearing in specific positions on a screen, and has to produce corresponding motor responses as quickly as possible. Given perceptual and motor limitations such as the need to shift visual attention to process a stimulus and to prepare a motor response before executing it, anticipating the location of the next stimulus and preparing the associated response allows for significantly faster reaction time. Learning in this task is measured by the improvement in response time between trial blocks in which the stimuli follow a repeated, deterministic sequence and those in which stimuli are

randomly selected. Our sequence-learning model (Wallach and Lebiere, 2000; Lebiere and Wallach, 2001) works by building representations (chunks) of small pieces of the stimulus sequence in working memory, storing them in long-term memory and retrieving them by matching to the most recent stimuli to predict the next item in the sequence. Perceptual-motor factors play an important role in this task as well, as the ability to learn the sequence and effectively use anticipation fundamentally depends upon the length of the interval between a response and the next stimulus. The model reproduces numerous behavioral measures, including average response times, probability of errors and percentage of anticipatory actions. Constraining models using multiple performance dimensions is essential to reducing degrees of freedom, a common problem in simple decision making tasks.

The essential feature of this model is its use of memories for specific experiences. Different experiences compete to be retrieved for use through an activation calculus that is based on the rational analysis of cognition (Anderson, 1990; Anderson and Schooler, 1991) of how the availability of memories is determined by the statistical structure of the environment. This fundamental idea is that the mind makes micro-decisions in retrieving an experience on which to base its next action. This insight has played out in a series of subsequent applications of this model to tasks that similarly involves making a sequence of decisions and performing associated actions.

The approach was then applied to a classic multi-person game, paper rock scissors (PRS). Games constitute an excellent decision making testbed because of the natural competitive pressure to make the best possible decisions and maximize performance. PRS is a two-person game in which each player has to simultaneously select one of three options. The winner is determined by a circular relation between the three options, with no option intrinsically better than the

others. Game theory prescribes random play as the optimal solution. However, while humans find it almost impossible to generate random actions, they find it quite natural to detect event sequences. Thus the sequence-learning approach is directly applicable to the iterated version of the game in which players engage in multiple rounds of play: the PRS model observes the opponent and learns small sequences of their moves in order to predict their next move, retrieves a best-matching sequence, and bases its move selection on that prediction. The model (Lebiere and West, 1999; West and Lebiere, 2001) matches quite well both aggregate level of human performance and specific characteristics such as the distribution of winning streaks. Its performance is also comparable to that of the best computer programs (Billings, 2000), an indication that cognitive constraints on decision making can provide useful functionality. The main aspect of this and other games for purposes of decision making is that the environment is not static but is instead another dynamic cognitive (human or model) entity adapting to one's actions, thus binding the players in a relation of reciprocal causation exhibiting signal detection characteristics such as stochastic resonance (West et al., 2005).

Sequential decision making is not confined to abstract games or experiments but instead is a natural component of many everyday situations. The game of baseball, specifically the competition between pitcher and batter, features the same structure of repeated choice among a set of possible actions. As in PRS, the pitcher has a number of options at his disposal varying in speed, location and movement, and, as in sequence-learning, because of perceptual-motor constraints the batter needs to anticipate the pitcher's choice in order to maximize the effectiveness of his response. Lebiere et al. (2003) applied the sequence-learning approach to two experimental situations. In the first one, pitch speed varied randomly between trials. The baseball model exploited the impact of recency in the base-level activation of the chunks representing each pitch to reflect the tendency of human batters to anticipate pitches similar to those they had seen recently. The fact that this pattern was observed in both humans and model despite the lack of any structure in the random trial-to-trial selec-

tion of pitch speed illustrates the pervasive nature of the biases inherited from the statistical structure of our environment. More fundamentally, it emphasizes that the concept of optimal decision making is relative to assumptions about the nature of the environment. For instance, probability matching, the common tendency to select choices in direct proportion to their quality, is often referred to as a suboptimal cognitive bias. However, that is only true if one assumes a fixed environment that one has adequately sampled. However, if one assumes a constantly changing environment, either independently or in response to our choices, probability matching can be an effective adaptive strategy to balance the need for constant sampling of the environment (exploration) with the goal to maximize performance given the currently available information (exploitation).

In the second experiment, pitch location was varied to reflect the current situation (specifically the balls and strikes count) to reflect strategic pitch selection in actual baseball games. The model's sensitivity to the context closely matched the strategic adaptivity of human batters. In both cases, the cognitive biases and mechanisms built into the architecture matched the human data *a priori* at least as well as a Hidden Markov model that had been trained on part of the data and could predict the rest *a posteriori*, emphasizing the role that cognitive constraints can play in modeling decision making biases in a principled, general basis rather than developing and parameterizing *ad hoc*, task-specific, models.

To demonstrate the relevance of the model to classical decision making paradigms and the power of cognitive constraints to *a priori* predict performance, Stewart et al. (2009) submitted a version of the PRS model to a choice prediction competition. The choice model uses the blending mechanism (Lebiere, 1999) to generate continuous expectations of the outcome of each option, reflecting both their payoffs and probabilities. The model won the part of the competition in which decisions between safe and risky options were based on prior experience with those options (Erev et al., 2010). Particularly remarkable is the fact that the sequence-learning model matched human performance better than models that attempted

to learn the value of each option in isolation despite, as in the first baseball experiment, the lack of any sequential structure in the task. But while the recency effect seems to be a pervasive bias ingrained in human cognition, not all such cognitive biases are equally stable and permanent. Lebiere et al. (2007) showed how the very same recency effect in the activation of experiences, weighted against a longer-term frequency effect, can explain the appearance and then disappearance of risk aversion bias in deciding between safe and risky choices. The recency effect initially dominates and leads selection away from risky choices, but the frequency effect then restores over time the balance between safe and risky options. This account is directly compatible with other sampling-based explanations (e.g., Denrell and Le Mens, 2007; Le Mens and Denrell, 2011) but provides precise predictions of the time course of the learning and unlearning of the risk aversion bias as a function of experience.

The direct implication of our account of decision making under uncertainty is the lack of stable risk preferences. Instead, the model attempts to achieve the best possible performance under cognitive and task constraints (such as payoff function and other performance metrics) without explicitly considering second-order information such as the amount of risk present in their decisions. In addition to the impact of learning and experience, cognitive factors influencing the level of risk assumed include individual differences parameters such as working memory capacity (Rehling et al., 2004) and noise in memory retrieval (West and Lebiere, 2001) as well as information framing effects (Martin et al., 2011).

## CONCLUSION

The application of the same model across a wide range of paradigms, from implicit learning of sequences to multi-person games in abstract and embodied settings to classical decision making tasks illustrates the predictive benefits of models based on cognitive architectures. Specifically, the cognitive constraints embedded in the architecture interact with the heuristic strategies used and the task environments to account for a broad pattern of results across multiple fields with limited parameter variations.

However, work remains to be done to achieve true unification. Instantiations of the model across paradigms still require the modeler to implement representation choices that reflect the nature of the task. A key part of decision making involves not only applying given heuristics and strategies, but also the metacognitive task of selecting among them and adopting the proper representation to implement them. Only when that aspect of decision making is viewed as an integral part of decision making and incorporated in models will a true theory of the field be achieved.

## ACKNOWLEDGMENTS

This research was supported by a Defense Threat Reduction Agency (DTRA) grant number HDTRA1-09-1-0053 to Christian Lebiere.

## REFERENCES

- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., and Lebiere, C. (2003). The Newell test for a theory of cognition. *Behav. Brain Sci.* 26, 587–637.
- Anderson, J. R., and Schooler, L. J. (1991). Reflections of the environment in memory. *Psychol. Sci.* 2, 396–408.
- Billings, D. (2000). The first international RoShamBo programming competition. *ICGA J.* 23, 42–50.
- Card, S. K., Moran, T. P., and Newell, A. (1983). *The Psychology of Human Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum.
- Curran, T., and Keele, S. W. (1993). Attentional and nonattentional forms of sequence learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 19, 189–202.
- Denrell, J., and Le Mens, G. (2007). Interdependent sampling and social influence. *Psychol. Rev.* 114, 398–422.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., Hertwig, R., Stewart, T., West, R., and Lebiere, C. (2010). A choice prediction competition, for choices from experience and from description. *J. Behav. Decis. Mak.* 23, 15–47.
- Gigerenzer, G., Todd, P. M., and the ABC Group. (1999). *Simple Heuristics that Make us Smart*. New York: Oxford University Press.
- Kruglanski, A. W., and Gigerenzer, G. (2011). Intuitive and deliberate judgments are based on common principles. *Psychol. Rev.* 118, 97–109.
- Kruglanski, A. W., Pierro, A., Mannetti, L., Erb, H., and Chun, W. Y. (2007). “On the parameters of human judgment,” in *Advances in Experimental Social Psychology*, Vol. 39, ed. M. P. Zanna (San Diego, CA: Zanna Elsevier Academic Press), 255–303.
- Le Mens, G., and Denrell, J. (2011). Rational learning and information sampling: on the ‘naivety’ assumption in sampling explanations of judgment biases. *Psychol. Rev.* 118, 379–392.
- Lebiere, C. (1999). The dynamics of cognitive arithmetic. Kognitionswissenschaft special issue on cognitive modelling and cognitive architectures, D. Wallach and H. A. Simon (eds). *J. German Cogn. Sci. Soc.* 8, 5–19.
- Lebiere, C., Gonzalez, C., and Martin, M. (2007). “Instance-based decision-making model of repeated binary choice,” in *Proceedings of the Eighth International Conference on Cognitive Modeling*, Ann Arbor, MI.
- Lebiere, C., Gray, R., Salvucci, D., and West, R. (2003). “Choice and learning under uncertainty: a case study in baseball batting,” in *Proceedings of the Twenty-fifth Annual Meeting of the Cognitive Science Society*, Boston, MA, 704–709.
- Lebiere, C., and Wallach, D. (2001). “Sequence learning in the ACT-R cognitive architecture: Empirical analysis of a hybrid model,” in *Sequence Learning: Paradigms, Algorithms, and Applications*, eds R. Sun and L. Giles (Germany: Springer LNCS/LNAI), 188–212.
- Lebiere, C., and West, R. L. (1999). “A dynamic ACT-R model of simple games,” in *Proceedings of the Twenty-first Conference of the Cognitive Science Society* (Mahwah, NJ: Erlbaum), 296–301.
- Martin, J. M., Juvina, I., Lebiere, C., and Gonzalez, C. (2011). “The effects of individual and context on aggression in repeated social interaction,” in *Proceedings of the HCI International 2011 Conference, HCII 2011*, Orlando, FL.
- Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Rehling, J., Lovett, M., Lebiere, C., Reder, L. M., and Demiral, B. (2004). “Modeling complex tasks: an individual difference approach,” in *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, August 4–7, Chicago, 1137–1142.
- Simon, H. (1957). “A behavioral model of rational choice,” in *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. New York: Wiley.
- Stewart, T. C., West, R., and Lebiere, C. (2009). “Applying cognitive architectures to decision making: how cognitive theory and the equivalence measure triumphed in the technion prediction tournament,” in *Proceedings of the Thirty-First Cognitive Science Conference*, Amsterdam.
- Wallach, D., and Lebiere, C. (2000). “Learning of event sequences: an architectural approach,” in *Proceedings of International Conference on Cognitive Modeling 2000* (NL: Universal Press), 271–279.
- West, R. L., and Lebiere, C. (2001). Simple games as dynamic, coupled systems: Randomness and other emergent properties. *J. Cogn. Syst. Res.* 1, 221–239.
- West, R. L., Stewart, T. C., Lebiere, C., and Chandrasekharan, S. (2005). “Stochastic resonance in human cognition: ACT-R vs. game theory, associative neural networks, recursive neural networks, q-learning, and humans,” in *Proceedings of the 27th Annual Conference of the Cognitive Science Society*, eds B. Bara, L. Barsalou and M. Bucciarelli (Mahwah, NJ: Lawrence Erlbaum Associates), 2353–2358.

Received: 16 August 2011; accepted: 11 October 2011; published online: 15 November 2011.

Citation: Lebiere C and Anderson JR (2011) Cognitive constraints on decision making under uncertainty. *Front. Psychology* 2:305. doi: 10.3389/fpsyg.2011.00305

This article was submitted to *Frontiers in Cognition*, a specialty of *Frontiers in Psychology*.

Copyright © 2011 Lebiere and Anderson. This is an open-access article subject to a non-exclusive license between the authors and Frontiers Media SA, which permits use, distribution and reproduction in other forums, provided the original authors and source are credited and other Frontiers conditions are complied with.