Available online at www.sciencedirect.com

SCIENCE DIRECT®

Cognitive Systems Research 6 (2005) 5–16

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Streak biases in decision making: data and a memory model

Action editor: Christian Schunn

Erik M. Altmann *, Bruce D. Burns *

Department of Psychology, Michigan State University, East Lansing, MI 48824, USA

Received 16 September 2004; accepted 16 September 2004

Abstract

Streaks of past outcomes, for example of gains or losses in the stock market, are one source of information for a decision maker trying to predict the next outcome in the series. We examine how prediction biases based on streaks change as a function of length of the current streak. Participants experienced a sequence of 150 flips of a simulated coin. On the first of a streak of heads, participants showed positive recency, meaning that they predicted heads for the next outcome with a greater-than-baseline probability. As streak length increased, positive recency first decreased but then increased again, producing a quadratic trend. We explain these results in terms of outcome-prediction processes that are sensitive to the historical frequency of streak lengths and that make heuristic assumptions about changes in bias of the outcome-generating process (here, the coin). An ACT-R simulation captures the quadratic trend in positive recency, as well as the baseline heads bias, in two experimental conditions with different coin biases. We discuss our memory-based model in relation to a model from the domain of economics that posits explicit representation of an "urn" from which events are sampled without replacement.

Keywords: Decision making; Memory; Cognitive simulation; ACT-R; Hot hand; Gambler's fallacy; Streaks

1. Introduction

A streak of repeated outcomes can be an important source of information for decision makers trying to predict the outcome of the next event. For example, asked to predict whether the global average temperature will increase or decrease next year, a decision maker who has access to historical records may predict an increase simply on the basis of past trends. A bias of this form, in which a streak of past outcomes is taken as evidence that the next outcome (or measurement) will be in kind, is often referred to as *positive recency*. In terms of

* Corresponding authors. Tel.: +1 517 353 4406; fax: +1 517 353 1652.
E-mail addresses: ema@msu.edu (E.M. Altmann), burnsbr@msu.edu (B.D. Burns).

1389-0417/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.cogsys.2004.09.002

the decision-maker's causal reasoning, one could imagine that a long streak of temperature increases induces a belief that an underlying causal mechanism is at work (warming due to greenhouse gases, for example). This causal mechanism will then govern changes in next year's measurement as well (modulo extraneous variance).

Of course, whether the decision maker makes this or the opposite prediction will depend on precisely how he or she represents the mechanism producing the streak. If the decision maker happens to work for the current Bush administration, he or she may well understand historical increases in the global temperature in terms more similar to the well-known gambler's fallacy, in which a gambler takes a streak of undesirable outcomes as evidence that his or her luck will soon change. More generally, a bias of this form, in which a streak of outcomes is taken as evidence that the next outcome will be opposite, is often referred to as *negative recency*. In the case of gambling, each gamble is an independent event, so there is no causal mechanism linking the outcomes (hence the fallacy). However, there are situations in which negative recency is a rational bias, namely when outcomes are sampled without replacement. For example, rats and other foraging animals have a bias against returning to the location where they found food on a previous trial (Olton, 1978). The underlying causal model is, presumably, that food at a given location is depleted before it is replaced; if this model is correct, then negative recency is adaptive.

In this paper we examine how recency biases change as a function of streak length. That is, we are interested in how the length of a streak, up to and including the most recent outcome, affects the decision maker's prediction concerning the next outcome. An experimental paradigm appropriate for addressing such issues involves a two-choice prediction task (similar studies complete with simulations include, e.g., Lebiere, Gray, Salvucci, & West, 2003; Lovett, 1998). In the standard experiment with this kind of paradigm, the participant is asked on each trial to predict the outcome of an event such as a coin flip. The "coin" typically has a bias toward one outcome or the other, of which the participant is not informed. The question of interest often has to do with probability learning – how the participant's bias to predict one outcome or the other changes over time. The usual finding is that participants "match" rather than "maximize", meaning that over many trials their bias tends to asymptote at the level of the bias in the event generator; for example, if the "coin" is biased to produce 75% heads, then participants will, by the end of a session, predict heads on roughly 75% of trials. Under a maximizing strategy, participants would come to predict heads 100% of the time, once they detected a bias, so a matching strategy is too difficult to explain using the simplest rules of rational choice.

Probability learning studies have thus shown that people are to some extent sensitive to base rates and changes in base rates, and adjust the frequency of their predictions accordingly, if suboptimally. Base rates, though related to streaks, are a distinct source of information, with different dynamics that may make them more or less appropriate to a given decision-making scenario. Thus, a probability learning experiment involving a biased coin might track changes in the bias to predict heads as experience with the biased coin grows. We are interested in tracking changes in the bias to predict heads as a streak of heads increases in length, from one head (following a tail), to two consecutive heads, and so on. Thus, in terms of the gambler's fallacy, we are interested in how the strength of the gambler's bias might change as a function of number of losses. Similarly, in terms of the hot-hand heuristic (Burns, 2004a), in which streaks of successes serve as an adaptive allocation cue, we are interested in how the strength of the team's bias to give the ball to one shooter is affected by that shooter's recent success at scoring.

Apart from the empirical question of whether decision makers respond to streaks, there is also an important theoretical question relating streaks to base rates. Both reflect any biases toward one outcome or the other in the outcome-generating process; the base rate of that outcome will be higher, and streaks of that outcome will be longer and more frequent. However, if the bias in the outcome-generating process happens to change, as a function of a shift in environmental characteristics,