Adaptable history biases in human perceptual decisions

Arman Abrahamyan[a,b,1], Laura Luz Silva[c,2], Steven C. Dakin[c,3], Matteo Carandini[c], and Justin L. Gardner[a,b]

[a]Department of Psychology and Neurosciences Institute, Stanford University, Stanford, CA 94305; [b]Laboratory for Human Systems Neuroscience, RIKEN Brain Science Institute, Wako-shi, Saitama 351-0198, Japan; and [c]UCL Institute of Ophthalmology, University College London, London EC1V 9EL, United Kingdom

Edited by J. Anthony Movshon, New York University, New York, NY, and approved May 9, 2016 (received for review September 21, 2015)

When making choices under conditions of perceptual uncertainty, past experience can play a vital role. However, it can also lead to biases that worsen decisions. Consistent with previous observations, we found that human choices are influenced by the success or failure of past choices even in a standard two-alternative detection task, where choice history is irrelevant. The typical bias was one that made the subject switch choices after a failure. These choice history biases led to poorer performance and were similar for observers in different countries. They were well captured by a simple logistic regression model that had been previously applied to describe psychophysical performance in mice. Such irrational biases seem at odds with the principles of reinforcement learning, which would predict exquisite adaptability to choice history. We therefore asked whether subjects could adapt their irrational biases following changes in trial order statistics. Adaptability was strong in the direction that confirmed a subject's default biases, but weaker in the opposite direction, so that existing biases could not be eradicated. We conclude that humans can adapt choice history biases, but cannot easily overcome existing biases even if irrational in the current context: adaptation is more sensitive to confirmatory than contradictory statistics.

decision making | choice bias | bias adaptation | choice history | computational modeling

When making decisions, we often learn from past failures and successes by using knowledge of those events to assist with subsequent choices. For example, when a choice leads to a reward, it can be beneficial to repeat it, or when it leads to failure, to avoid it. Such a success-stay/fail-switch strategy can be fruitful in cooperative behavior (1). Indeed the literature on reinforcement learning is built on the idea that we learn the value of choice options from the outcomes of past decisions (2, 3). Moreover, when the only source of information on expected value are past decisions and their outcomes, subjects have little difficulty relying on these past cues (4). These behaviors demonstrate that humans can appropriately adapt to the statistics of their history of choices.

Humans and other animals, however, also apply strategies based on past failures and successes in contexts where their use is irrational and will adversely affect performance. For example, human subjects tend to apply the success-stay/fail-switch strategy to the game of "rock-scissors-paper" (5) as do monkeys in the "matching pennies" task (6). The optimal behavior, instead, is to respond randomly, because any strategy can be exploited by other players (7). People also resort to the suboptimal success-stay/fail-switch strategy in situations that would require more complex decision-making strategies (8). Thus, despite our ability in some situations to adapt to choice history statistics, in other contexts, we are unable to appropriately adjust behavior to account for choice history statistics.

Sensory psychophysics offers an opportunity to examine how subjects integrate sensory evidence with past history and what conditions promote appropriate adaptation to choice history statistics. Most psychophysics experiments are carefully designed such that subjects should make choices based only on the sensory evidence at hand. However, under such conditions, human observers exhibit nonsensory biases (9–20) For instance, in the classical two-alternative forced choice task, a subject is asked to make a perceptual decision based solely on the present sensory evidence (21) and to use success or failure feedback only to optimize their use of such sensory information (22). However, both mice (14) and humans (15, 17, 18) make choices that are biased by their recent history of successes and failures and are thus suboptimal. Such biases, where prior beliefs rather than evidence guide behavior, are a form of fallacy, similar to the gambler's fallacy or hot-hand fallacy (23, cf. 24).

Why might choice history biases remain despite degrading performance? One possibility is that they may be hard-wired, inadaptable to the statistics of the task at hand. In support of this view, an early study reported that choice history biases remained constant regardless of whether stimulus sequences were ordered or randomized (25). Perhaps choice history biases are not adaptable or perhaps they cause such a small loss in performance that there is not sufficient incentive to adapt them (6) as they may be optimal in more general circumstances.

To address this question, we assessed whether choice history biases can adapt when given a learning signal in the form of trial order statistics. We manipulated the probability of a visual stimulus appearing on the left or right side of the screen, depending on the success or failure on the previous trial. For example, in one of the conditions, if the subject failed on the previous trial by choosing left, the stimulus would be presented on the right with 80% probability. We used a probabilistic choice model (14) to document each subject's choice history biases, such as preference to stay on the same side after a failure or switch after a success. By measuring the biases, this model allowed us to assess their adaptability, by comparing the conditions where trial statistics were manipulated to the baseline condition where they were truly random. We found that

Significance

Adapting to the environment requires using feedback about previous decisions to make better future decisions. Sometimes, however, the past is not informative and taking it into consideration leads to worse decisions. In psychophysical experiments, for instance, humans use past feedback when they should ignore it and thus make worse decisions. Those choice history biases persist even in disadvantageous contexts. To test this persistence, we adjusted trial sequence statistics. Subjects adapted strongly when the statistics confirmed their biases, but much less in the opposite direction; existing biases could not be eradicated. Thus, even in our simplest sensory decisions, we exhibit a form of confirmation bias in which existing choice history strategies are easier to reinforce than to relinquish.

Author contributions: A.A., S.C.D., M.C., and J.L.G. designed research; A.A., L.L.S., and M.C. performed research; A.A., L.L.S., and M.C. analyzed data; A.A., M.C., and J.L.G. wrote the paper; and S.C.D. provided software and hardware and organized experiments.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: armana@stanford.edu.

[2]Present address: UCL Philosophy Department, University College London, London WC1E 6BT, United Kingdom.

[3]Present address: School of Optometry and Vision Science, University of Auckland, Auckland 1142, New Zealand.

choice history biases are highly adaptable. However, reversing an existing bias to switch after failure proved more resistant to adaptation than promoting existing biases.

We presented preliminary versions of these results in two distinct conference abstracts (15, 17). Given the concordance of the results, we decided to join forces and publish the work together.

Results

To estimate choice history biases, we performed classical measurements of contrast sensitivity using a two-alternative forced choice design (Fig. 1 *A* and *B*). We presented a visual grating on the left or the right of the screen, as asked the subjects to report its position by pressing one of two buttons. We varied grating contrast to make the task easier or harder. We collected data in three laboratories situated in Japan, the United Kingdom, and the United States. At RIKEN and Stanford, the grating appeared on a gray background (Fig. 1*A*). At University College London (UCL), the grating appeared on one of two patches of visual noise (Fig. 1*B*). Subjects were informed of their success or failure by a sound (at RIKEN and Stanford) or a change in color of the fixation cross (at UCL). These methods are standard in visual psychophysics and have been used for over a century (26–28).

Identifying Choice History Biases. In a first set of experiments, we fully randomized the contrast and position of the stimuli, so that past trial success or failure had no influence on the current trial. Consequently, adopting any choice history strategies by subjects was suboptimal for this task. Nevertheless, we found that subjects in all three laboratories showed clear choice history biases.

To gain initial insight into these biases, we computed the proportion of rightward responses (psychometric function) for each subject after sorting trials into two sets: trials preceded by a left choice and trials preceded by a right choice (Fig. 1 *C*–*E*). By computing proportion of rightward choices separately for those two sets, we could visually assess if the subject had a tendency to switch or to stay. For some subjects, there was an equal number of staying and switching, suggesting they did not have choice history biases (Fig. 1*C*). However, other subjects tended to switch more often than staying (Fig. 1*D*). Yet other subjects preferred to stay rather than to switch (Fig. 1*E*). Overall, switching/staying biases were seen when stimulus intensities were harder to perceive (the middle part of the psychometric curves), suggesting that subjects were defaulting to choice history biases when sensory evidence was weak.

Quantifying Choice History Biases. To describe and quantify these biases, we fit a probabilistic choice model to the responses (14) (Fig. 2*A*). In the model, a choice is made by flipping a coin with unequal odds. The logarithm of the odds is the sum of a sensory term (that depends on current stimulus side and contrast) with two choice history terms (which weigh previous successes and failures) and with a general bias for left or right. If, for example, a subject displayed a tendency to switch sides after a failure, the failure weight would be negative. If the subject tended to stay on the same side, the failure weight would be positive. If failure had no effect, its weight would be zero.

This scheme could also capture a general tendency to switch or stay irrespective of success or failure. For a subject who tends to switch regardless of outcome, both success and failure weights would be negative. Conversely, for a subject who tends to stay regardless of outcome, both weights would be positive. The impact of success or failure on a subject's decision, therefore, can be assessed by asking whether the weights associated to success and failure are statistically different. Similarly, the impact of choice history on a subject's decisions can be assessed by asking whether the weights associated with past choices are significantly different from zero.

Most subjects were significantly biased by their choice history (Fig. 2*B*). We used a likelihood ratio test to compare the full model to a "no-history" model where choice history weights were

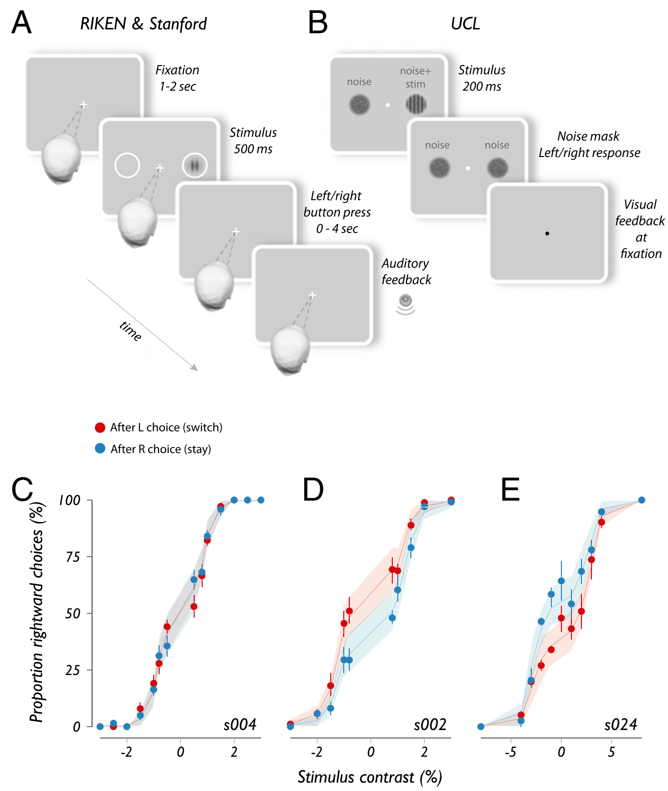

Fig. 1. Task design and examples of psychometric functions biased by the previous choice. Subjects performed a two-alternative forced-choice discrimination of whether a sinusoidal grating presented against a gray background (*A*; RIKEN and Stanford subjects) or superimposed on visual noise (*B*; UCL) was to the left or right of fixation. (*C*–*E*) Examples of psychometric curves for three subjects sorted by whether the previous trial was a left (red) or right (blue) choice. Stimulus contrast (abscissa) is coded as positive for a stimulus on the right and negative for the left. Examples of subjects without any bias (*C*), a tendency to switch sides (*D*), and a tendency to stay (*E*) are depicted. In each case, a probabilistic choice model fitted to each subject's data accurately fitted these effects (lines and shaded areas indicating 68% CIs). Error bars for the data are bootstrapped SEM.

set to zero. Applying this test for each run (each subject completed multiple runs), we found that the full model better explained each subject's choices for the vast majority of runs [66 ± 6.2% SE of runs had $\chi^2(2) \leq 0.05$ uncorrected for multiple comparisons; blue bars in Fig. 2*A*]. In terms of subjects, the full model better explained at least 75% of the runs for 18 subjects (i.e., 50% of subjects). By this measure, only data from seven subjects (19%) could be explained by the no-history model. We also confirmed the results of the likelihood ratio test using leave-one-trial-out cross-validation on each run and found that the majority of runs had a better median likelihood for the full model than for the no-history model on the left-out trial. If anything, the cross-validation analysis found 9% more runs better fit with the full-model than the likelihood ratio analysis.

The importance of the choice history terms can also be judged from the model's predictions of the psychometric curves (Fig. 1 *C*–*E*). The model did not assign large choice history terms to the first example subject ($\beta_{Success} = 0.01$, $\beta_{Fail} = -0.17$; Fig. 1*C*), but it did for the subsequent two subjects. The subject with a tendency to switch sides had a large failure bias ($\beta_{Fail} = -0.89$; Fig. 1*D*), and the subject with a tendency to stay had a large success bias ($\beta_{Success} = 0.60$; Fig. 1*E*). As shown by the confidence intervals (CIs), with these parameters the model provided excellent fits to the data.

Choice history biases had a sizeable effect on perceptual decisions, an effect equivalent to that elicited by a low contrast stimulus

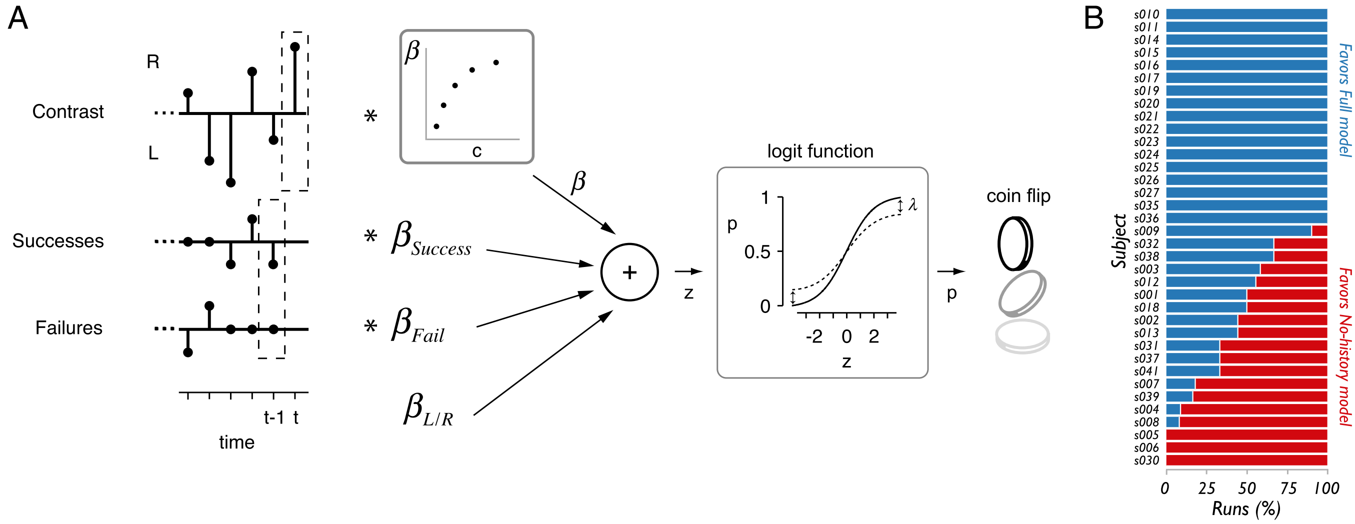

Fig. 2. Probabilistic choice model and statistical test for the role of choice history terms. (*A*) The probabilistic choice model represents choices as a linear sum of sensory evidence (contrast), choice history biases (successes and failure), and general L/R bias as predictors. Dashed boxes show example of predictors from one trial. Fitted weights of the model provide the estimate of the magnitude of influence of sensory and nonsensory terms rectified using the lapse rate (λ). The weighted sum can be transformed into choice probability using the logistic function. The model can then simulate trial-by-trial choices by "flipping the coin" using choice probability *p*. Modified with permission from ref. 14. (*B*) Proportion of runs for each subject for which the full model (blue) or no history model (red) provided better fits according to a likelihood ratio test.

(Fig. 3 *A–C*). As expected, subjects were more influenced by the stimulus when it had high contrast: the weights associated with the stimulus invariably grew with stimulus contrast (Fig. 3 *A–C*). Overall contrast intensities (abscissa of Fig. 3 *A–C*) differed between measurements made at RIKEN and Stanford and those made at UCL due to differences in the stimuli (*Materials and Methods*). Choice history biases, however, did not depend on specifics of the stimulus or the testing laboratory. The weights associated with these biases were on average smaller than those associated with high-contrast stimuli but of similar magnitude as those for lower contrasts. The weights associated with failures, specifically, tended to be negative and of magnitude nearing those associated with contrasts close to the average 75% detection threshold. Conversely, the average success bias was positive. These results indicate that subjects on average used a success-stay/fail-switch strategy.

We verified that history biases were not different across three testing locations [Japan, United States, or United Kingdom; $F(2,33) = 0.49$, $P = 0.62$, partial $\eta^2 = 0.03$]. However, history biases could differ substantially across subjects (Fig. 3*D*). Some subjects showed a strong bias to switch sides after a failure (negative failure weights on lower portion of vertical dashed line; Fig. 3*D*), and some tended to stay after a success (positive success weight on right portion of horizontal dashed line; Fig. 3*D*). A few, as noted above, did not express any significant choice history biases.

A possible post hoc hypothesis for this diversity is that subjects better trained and more knowledgeable about psychophysical tasks would have less of a tendency to have history biases, i.e.,

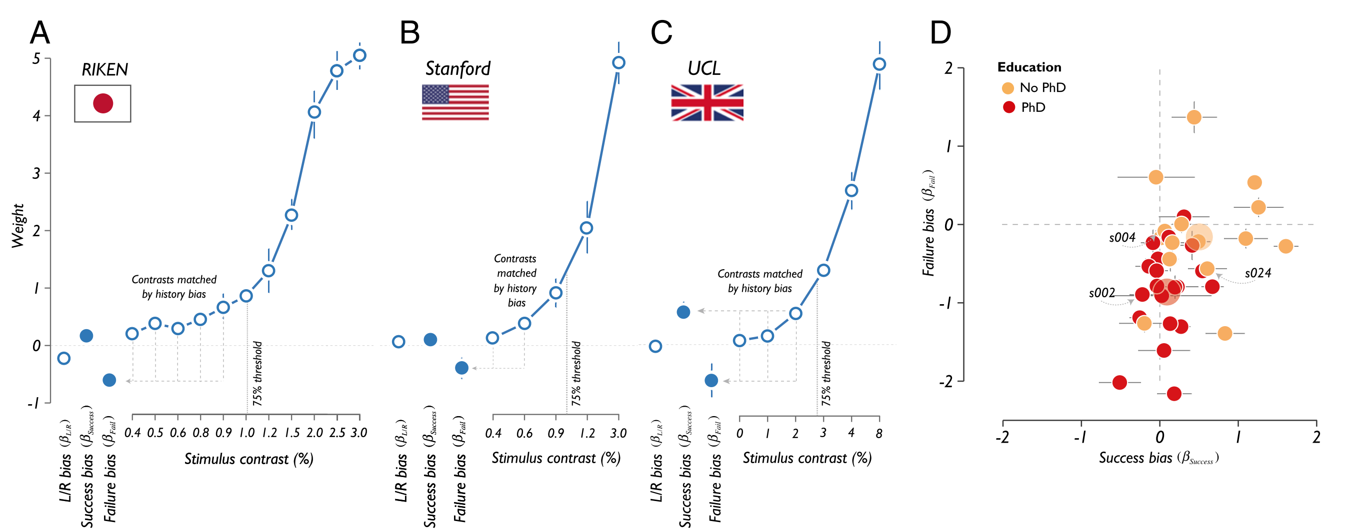

Fig. 3. Quantifying choice history biases. (*A–C*) Choice-history and contrast weights of probabilistic choice model averaged across subjects for data collected across diverse demographics at RIKEN (*A*), Stanford (*B*), and UCL (*C*). Error bars are bootstrapped SEM. (*D*) Success and failure biases of individual subjects colored according to whether subjects had (or were in the process of obtaining) a PhD (red, large dot is mean across these subjects) or not (orange). Example subjects from previous figures are indicated. Error bars are SEM.

nonoptimally biased by choice histories. As our subjects came from diverse backgrounds, we examined choice history biases using graduate education as a proxy measure for experience and knowledge of psychophysical tasks. We split our subjects by education into a PhD group (those with or working toward a PhD, $n = 20$) and no PhD group ($n = 16$).

We found that education level did significantly covary with choice history biases. However, contrary to the hypothesis, both types of subjects showed significant irrational history-based biases: PhD subjects were more likely to be affected by failure and the others more likely to be affected by success. Specifically, PhD subjects tended to have a "switch after failure" bias [red points, Fig. 3*D*; mean $\beta_{Fail} = -0.86$, $t(19) = 6.5$, $P < 0.01$; mean $\beta_{Success} = 0.09$, $t(19) = 1.5$, $P = 0.33$; *t* tests comparing to 0]. By contrast, no PhD subjects tended toward a "stay after a success" bias [yellow points; mean $\beta_{Success} = 0.50$, $t(15) = 3.70$, $P < 0.01$; mean $\beta_{Fail} = -0.16$, $t(15) = 0.95$, $P = 0.36$; *t* tests comparing to 0). *P* values were Bonferroni adjusted for multiple comparisons.

We stress that these results, however, are post hoc observations on subject samples that were not equated for potentially important factors such as age or IQ and do not exclude the possibility of other covariates such as sex, ethnicity, age, or even whether subjects majored in Psychology or Neuroscience. All these factors could affect the results. We report them here solely as a potential starting point for future investigation into covariates of choice history biases.

Further supporting the view that history biases were not due to inexperience, we found that biases did not decrease across task runs. We regressed success and failure bias weights against run number and found no significant effects ($slope_{Success} = 0.03$, $slope_{Fail} = -0.01$, both $P > 0.05$, regression to individual z-score values of success and failure biases). These results indicate that subjects commenced the task with preexisting biases and did not learn to adapt them over the course of a few thousand trials.

Quantifying History Bias-Driven Sensitivity Loss. Using model simulations, we found that the history biases caused a significant loss in visual sensitivity (Fig. 4). We measured sensitivity loss as the ratio of slopes of psychometric curves computed with and without bias (Fig. 4*A*). We found that 67% of subjects, or 24 of 36, had significant sensitivity loss due to choice history biases (all $P < 0.01$, Wilcoxon one-sample median test comparison with 0). The mean loss of sensitivity was 4.7% (95% CI, 3–7%). For 19% of subjects, the bias-driven decline ranged from 10% to 20%. To visualize the sensitivity loss across subjects we constructed a model based on the average sensory weights from all subjects, and we systematically varied history biases in a range of failure and success weights similar to that found in individual subjects (Fig. 4*B*).

Inducing Choice History Biases. Having found that subjects exhibited consistent and significant choice history biases even when these biases worsen performance and decrease sensitivity, we asked if these biases are immutable or whether they are adaptable. If subjects can adapt choice history biases given a large enough incentive, then we should be able to induce choice history biases by adjusting trial statistics. In a subset of subjects from RIKEN and Stanford, we adjusted trial statistics such that 80% of the time after a specific outcome (success or failure), the target would have a predictable position (stay or switch sides).

The results revealed that both failure and success biases were highly adaptable (Fig. 5). When we changed the trial statistics after failure, we found a significant shift in failure bias in the optimal direction (Fig. 5*A*). The subjects adopted a more negative failure bias (making them more likely to switch from a losing side) when failure led the targets to switch than when it led them to stay [mean difference, $\beta_{Fail} = 0.45$, $t(13) = 3.00$, $P = 0.02$; 95% CI, 0.13–0.78; Bonferroni-adjusted]. Conversely, as may be expected from optimal behavior, the success bias was not significantly affected [mean difference of $\beta_{Success} = 0.06$, $t(13) = 0.46$, $P = 0.99$; 95% CI, −0.32 to 0.21; Bonferroni-adjusted]. Similarly, when we changed the trial statistics after success, we found a significant shift in success bias in the optimal direction (Fig. 5*B*). The subjects adopted a more positive success bias (making them more likely to stay on a winning side) when success led the target to remain in the same position than when it led it to switch sides [Fig. 5*B*; mean difference of $\beta_{Success} = 1.04$, $t(13) = 6.61$, $P < 0.01$; 95% CI, 0.70–1.38; Bonferroni-adjusted]. The subjects also slightly adjusted their failure bias [mean difference

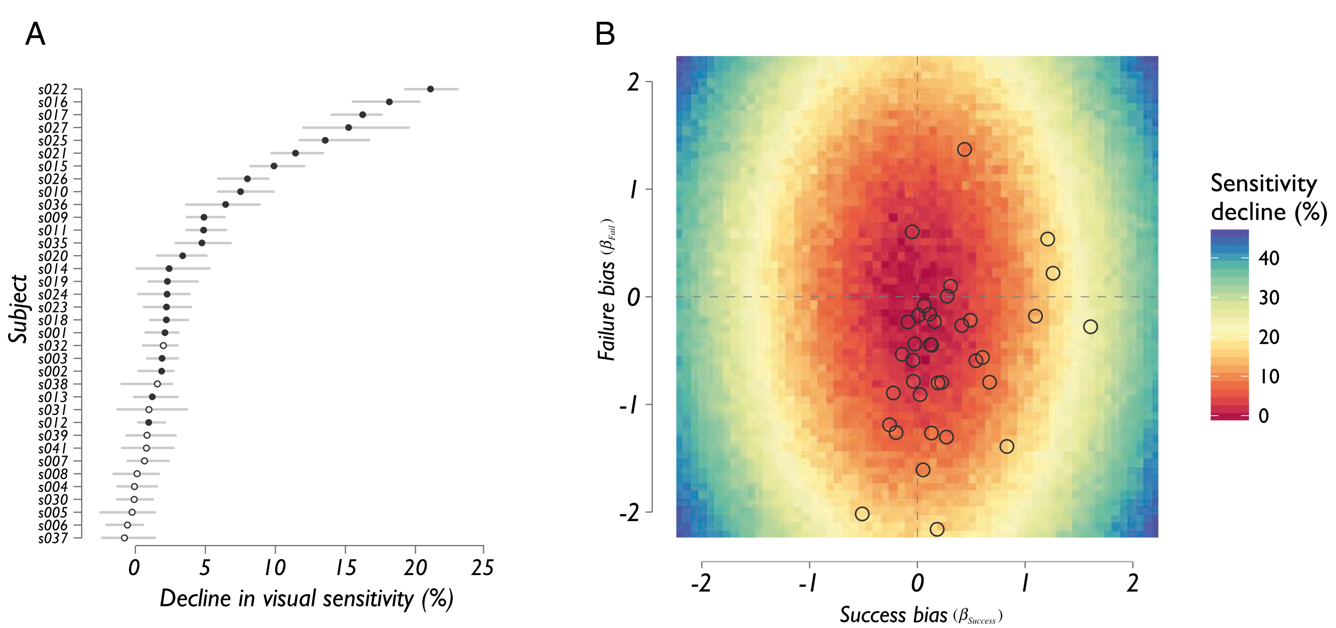

Fig. 4. Bias-driven sensitivity loss. (*A*) Biases can significantly reduce visual sensitivity when comparing the slope of psychometric function from responses simulated with and without biases. Black circles show subjects with significant median decline in visual sensitivity ($P < 0.01$). Error bars are median bootstrapped 95% CI. (*B*) Sensitivity decline matrix shows simulated median loss in sensitivity as a function of choice history biases using average sensory weights across all subjects. Each grid point shows the median decline in sensitivity over 200 simulation runs. Circles show individual mean biases.

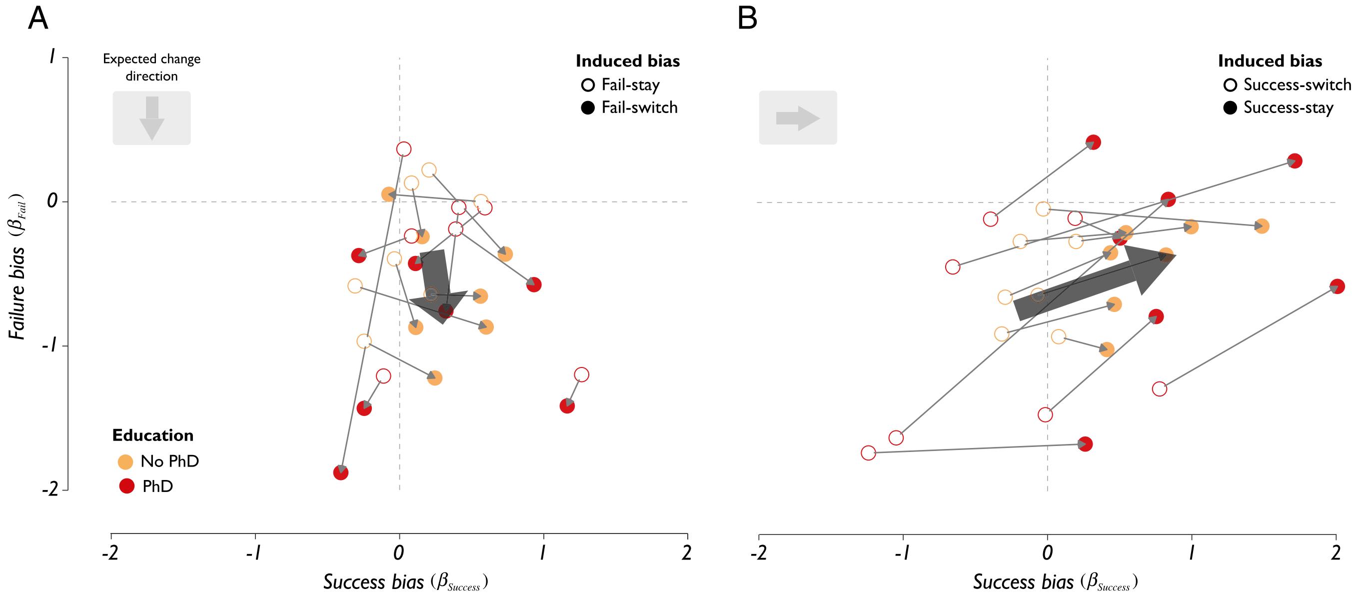

Fig. 5. Induced choice history biases. (*A*) Subject by subject choice history weights for experiments in which trial statistics were manipulated such that on 80% of trials stimulus presentation location was switched after a failure (○) or stayed on the same side (●). The gray arrow shows the mean group shift from origin to displacement. (*B*) Same conventions as *A*, but for when trial statistics were manipulated after successes.

of β_{Fail} = 0.36, $t(13)$ = 2.76, P = 0.03; 95% CI, 0.07–0.63; Bonferroni-adjusted], which suggests a bit of overgeneralization.

Although both bias weights were adaptable, induced failure biases were on average half as large as induced success biases. We reasoned that the difference in adaptability between success and failure biases could arise from the difference in these biases already observed in the control condition, when successive trials were independent (Fig. 3*D*). Indeed, in many subjects, the failure bias was substantially stronger than the success bias. Perhaps a strong pre-existing bias is less adaptable than a weak one?

To test this hypothesis, we compared the magnitude of the shift of induced bias relative to subjects' natural bias (Fig. 6). We used the same data as reported above, but compared each condition to the control condition, in which trial statistics were completely random. As we have seen (Fig. 3*D*), in the control condition subjects tended to have negative failure biases (leading to a tendency to switch after failure) and weaker, positive success biases (leading to a weaker tendency to stay-after-success, mean β_{Fail} = –0.55, mean $\beta_{Success}$ = 0.27, both P < 0.01, Bonferroni adjusted). We take these values to indicate each subject's "natural bias" (open symbols in Fig. 6). We then examined the four adaptation conditions one by one.

This analysis confirmed that adaptation was more successful in shifting biases in the same direction as a subject's natural biases. Specifically, when the stimulus statistics favored an unnatural strategy of staying after failure (Fig. 6*B*), adaptation of the failure bias was small and not significant [mean difference of β_{Fail} = 0.06, $t(15)$ = 0.59, P = 0.99; 95% CI, –0.29 to 0.17; Bonferroni adjusted]. By contrast, when the stimulus statistics encouraged the natural switch-after-failure bias (Fig. 6*A*), it resulted in significant effects in the expected direction [mean difference of β_{Fail} = 0.39, $t(14)$ = 5.4, P < 0.01; 95% CI, 0.23–0.54; Bonferroni adjusted). Similarly, subjects significantly changed their success bias when stimulus statistics favored the natural, but weak, tendency of staying after a success [Fig. 6*C*; mean difference of $\beta_{Success}$ = 0.64, $t(16)$ = 5.66, P < 0.01; 95% CI, 0.40–0.88; Bonferroni adjusted]. Subjects also slightly adapted to the opposite unnatural strategy of switching after a success [Fig. 6*D*; mean difference of $\beta_{Success}$ = 0.36, $t(13)$ = 2.90, P = 0.03; 95% CI, 0.09–0.63; Bonferroni adjusted).

Although there was a large bias variability between subjects in the adaptation condition, the results were not driven by subjects who changed their biases the most. We verified the absence of long tails by computing D'Agostino's test of normality and skewness for both failure and success induced bias in each condition, and we could not reject the hypothesis that weights were normally distributed and not skewed (all P > 0.05).

To examine whether induced history biases increased across runs, we calculated the slope of a linear regression fitted to success and failure biases against runs in a given adaptation condition (subject z-score values computed by bias type and bias induction condition). Subjects gradually increased their stay-after-success bias ($slope_{Success}$ = 0.21, P < 0.01; no change in failure bias) and their switch-after-success bias ($slope_{Success}$ = –0.41, P < 0.01; no change in failure bias). Failure biases, instead, did not significantly change across runs (all P > 0.05).

One possibility is that subjects could use a conscious cognitive strategy if they became aware of our trial history manipulations. However, when debriefed at the end of the experiment, no subject noticed that switching or staying was contingent on previous choices, despite the fact that their history weights showed that they had learned this contingency.

Finally, we examined whether subjects were better able to adapt history biases after being exposed to our trial history manipulations. To examine changes in adaptability to trial history manipulations, we split the data into "early" vs. "late" runs and compared the magnitude of adaptation between the two. This comparison is fair, because the order of adaptation conditions was randomized and counterbalanced to avoid block ordering effects. When subjects adapted to success-stay bias late rather than early, they did show larger success biases [$\beta_{Success\_early}$ = 1.1, $\beta_{Success\_late}$ = 0.4, $t(12.9)$ = 3.8, P < 0.01, Bonferroni adjusted]; however, no other condition showed this effect (all P > 0.05). Overall, therefore, there seems to be an innate ability to adapt the history biases to the trial statistics, especially when these statistics tend to encourage the subject's default biases.

Discussion

We have seen that subjects display systematic choice history biases in their behavior, even when performing a psychophysical task where trials are fully randomized, and such biases lead to poorer performance. These biases generally manifest as switch-after-failure

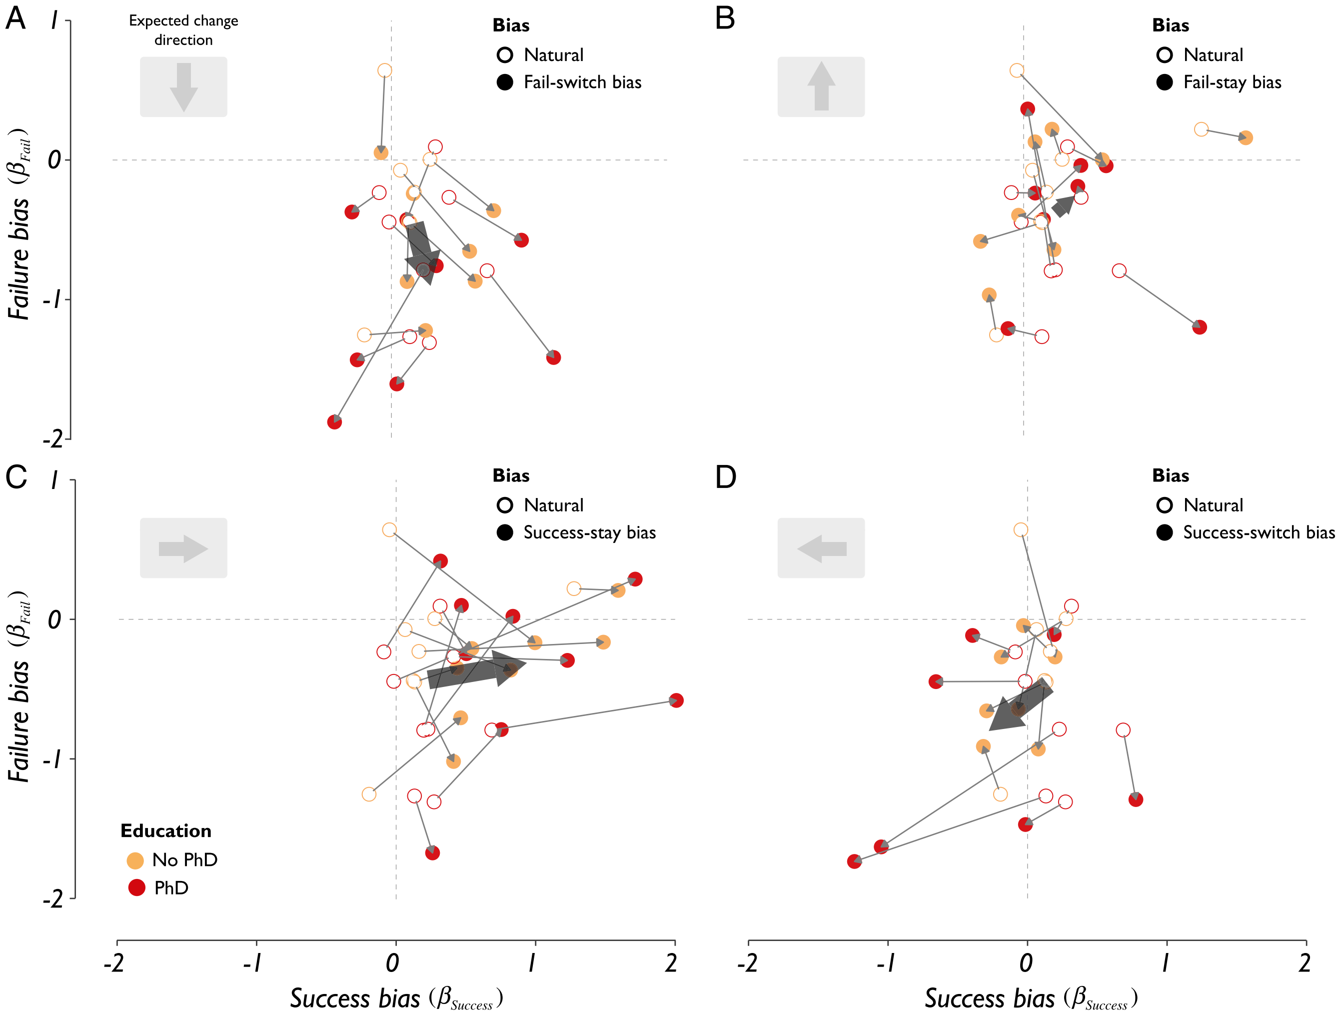

Fig. 6. Biases are easier to induce when they align with subject's natural biases. Data from Fig. 5 are plotted for each condition against experimental runs in which trial order was completely randomized (○), which we use as a measure of subject's natural bias. Subjects had a tendency for switch-after-failure and stay-after-success (gray arrows all begin in lower right quadrant). Inducing in this same direction (*A* and *C*) resulted in large adaptation effects, whereas inducing in the opposite direction (*B* and *D*) resulted in small if any changes. Gray arrows show the mean group shifts from origin to displacement. Plotting conventions similar to Fig. 5.

or as stay-after-success strategies (8, 10, 11, 18). Both types of biases were disadvantageous in that they reduced sensitivity for most of the subjects. By manipulating the statistics of the trials so as to encourage or discourage these strategies, we found that the weights were most adaptable in a direction that strengthened them, and much less so in a direction that reduced them. Our results thus demonstrate that subjects can alter, but never fully disengage, their natural choice history biases.

Although our results show a difference in adaptability, they do not prove inadaptability of existing strategies. Perhaps with stronger incentives, even such existing strategies might be adaptable. History biases thus act like an unconscious confirmation bias; when statistics of the environment agree with existing biases, they are strengthened, but they are resilient to nonconfirmative statistics.

The inability to disengage disadvantageous choice history biases suggests a more global strategy; indeed such biases occur across many types of stimuli and sensory modalities. Choice-history biases have been observed when subjects had to make judgments about physical weights (9), auditory stimuli (12, 16, 18), or visual stimuli (4, 8, 10, 11, 14, 18). As in our experiment, subjects show a diversity of biases ranging from switching strategies (4, 9, 12, 14, 16), to staying strategies (10, 11, 18), to success-stay/fail-switch strategies (5, 6, 8, 18). These biases are typically limited to the preceding one trial and the magnitude of the choice history biases is inversely proportional to the strength of sensory stimulus, such that weaker sensory stimulus elicits stronger choice history biases (13, 15, 18). These effects are readily captured by our simple choice history model (Fig. 2). However, it is possible that more complex choice history biases exist, which may be identified using information theory or novel statistical methods (29).

Our results and our model provide a way to correct for the loss in visual sensitivity caused by choice history biases. This loss was statistically significant for most subjects with a mean near 5%. For about 20% of subjects, it resulted in a loss that ranged from 10% to 20%. Correcting for these kinds of biases can have implications not only for basic but also for clinical and applied vision science. Visual sensitivity is typically measured using a two-alternative force choice (2AFC) task similar to the one we used here (30). Coupled with the fact that older adults maybe more prone to using choice history biases than younger adults (31), correcting for choice history bias driven visual threshold measurement errors may be important in real-world applications of vision sensitivity testing.

That choice history biases were most prevalent at weaker stimulus intensities is suggestive of a Bayesian inference strategy, where biases act as priors that influence decisions more when sensory evidence is weak (32). Given this, Bayesian models of contrast discrimination (33–35) might be better extended to incorporate choice history biases than models that do not incorporate priors (27, 36). Bayesian frameworks have been used to explain motion perception biases (37, 38), biases around cardinal orientations (39) and tilt perception biases (40). An ecological basis of these biases is thought to come from long-term and evolutionary exposure to environmental statistics (37, 39), although stimuli seen a few seconds earlier can also bias our judgment of orientation perception (41, 42). Thus, perceptual biases are thought to be specific to the particular statistics of the sensory quality that is being inferred. Like what we found with choice history biases, training can also alter more specific perceptual biases (43). Similarly, stimulus expectations can be introduced by cueing subjects to probabilities of motion direction or indicating payoffs associated with one or another choice (44–46). Despite these similarities, perceptual biases and perceptual learning (47) effects typically do not readily transfer to other stimuli and tasks, and are thus likely different from choice history biases, which tend to generalize across stimuli and tasks.

If choice history biases reflect more global strategies than those in perceptual inference and then reinforcement learning (2), with some crucial caveats, might provide another computational framework for understanding them. Reinforcement learning provides a compelling account of how animals and humans can learn the value of choice options based on past choices and rewards. In games where reinforcement learning has been used to model behavior, animal (48, 49), and human (8) subjects often display success-stay/fail-switch strategy similar to those observed in our experiments. Our task differs from typical reinforcement learning tasks in that it requires weighing sensory evidence and choice options where the latter have no explicit value or probability associated with them (50). Nonetheless, humans may perform the task by assigning value to choice options (51) or by treating successful trials as rewards. Unlike typical reinforcement learning tasks in which the value of a perceptually strong stimulus is assigned a value, choice history biases were most evident at weak contrasts, suggesting that they may best be described by an amalgam of Bayesian inference and reinforcement learning theories.

Taken together, our results suggest that choice history biases are general, not easily overcome in individual tasks and thus a strategy likely to be useful across a wide variety of different contexts. Choice history biases are evident in a variety of situations even when they are maladaptive (5). A possibility is that subjects are applying strategies that lead to optimal decisions in their natural environment, but those same strategies become fallacies when faced with the artificial constrains of psychophysical experiments (52, 53). In that sense, they resemble heuristics, which are commonly used to assist in decision making under conditions of uncertainty (54). Heuristics are generalizable and can be flexibly used in many decision-making contexts (55), but when exposed in a context where they are maladaptive, they seem arbitrary and problematic. Similarly, choice history biases appear as fallacies in attempts to rationally measure sensation, but nonetheless likely constitute an important component of choice behavior across a wide array of environments.

Materials and Methods

Observers. Fourteen naive observers and one of the authors volunteered as subjects (five females; mean age, 31 y; age range, 21–38 y) for data collected at RIKEN. At Stanford University, 12 subjects volunteered, but 3 were excluded as a result of exceptionally poor performance that suggested they were not preforming the task (leaving 7 females and 2 males; age range, 19–36 y). At the UCL, data from 12 subjects were collected (1 female; mean age, 35 y). Subjects gave prior written informed consent and optically corrected their vision when necessary. The study procedures were approved in advance by the RIKEN Ethics Committee and local Ethics Committees at the UCL and Stanford.

Apparatus at RIKEN and Stanford. Sitting in a dark room, observers responded to stimuli presented on a 21-in gamma-linearized flatscreen cathode ray tube (CRT) monitor (Dell Trinitron P1130; Dell) at RIKEN and 22.5-in light-emitting diode (LED) monitor at Stanford both operating at 100-Hz vertical refresh rate with resolutions set at 1,980 × 960 or 1,920 × 1,080 pixels, respectively. Observers self-adjusted the height of a pneumatic chair or table to comfortably place their head on a gel-cushioned chin and forehead rest. The chin and forehead rest restricted head movements which facilitated accurate monocular eye-tracking using infrared video-based eye-tracker at 500 Hz (EyeLink 1000; SR Research). We calibrated the eye tracker at the beginning of each run using a built-in five-point calibration procedure. Visual stimuli and task sequence were programmed in MATLAB (MathWorks) using the MGL library (justingardner.net/mgl) on a Mac Pro computer (Apple).

Apparatus at UCL. Subjects responded to stimuli presented on a 21-in CRT monitor operating at 75 Hz and gamma-linearized in software. Stimuli and trial sequences were generated using a Matlab script and Psychtoolbox (56, 57).

Task and Stimulus.

RIKEN and Stanford. Subjects were asked to detect faint visual stimuli that were presented for 500 ms either to the left or right of fixation. Stimuli were vertical sinusoidal gratings (1 c/°, 6° in width and height, modulated by a symmetric 2D Gaussian window at 100% of peak contrast) presented 12° to the left or right of fixation at a viewing distance of 50.5 cm at RIKEN and 56 cm at Stanford (Fig. 1*A*). Gratings randomly drifted either to the left or to the right at 0.5 c/s. The stimulus was presented within a white ring (diameter, 7° of visual angle) to reduce spatial uncertainty (36). The monitor was calibrated with a Topcon SR-3A-L1 Spectroradiometer (Topcon) at RIKEN and SpectraScan 650 (Photo Research) at Stanford. The gamma table was linearized and dynamically adjusted to take full advantage of the 10-bit gamma table resolution. The background appeared uniformly gray with luminance equal to 47 cd/m^2 (midpoint of calibrated monitor's full luminance range) both at RIKEN and Stanford.

Stimulus detection difficulty varied with stimulus intensities that ranged from 0.4% (very difficult) to 3% (very easy) Michelson contrast, which is computed by taking the difference between the lowest and highest luminance values of the stimulus and dividing by the full luminance range of the monitor. We took advantage of the video card's 10-bit gamma table to improve luminance resolution while maintaining luminance linearity. Subjects did not report image aftereffects.

At the beginning of each trial subjects fixated at the central white cross for 1–2 s followed by the presentation of the stimulus either to the left or right side of the fixation. When the stimulus disappeared, the fixation color changed to pale blue prompting subjects to respond. The task instruction was to fixate, detect the stimulus location and rapidly and accurately make a choice by pressing one of two keys with their left hand. An auditory feedback followed, which indicated correct and incorrect decisions ("pop" and "basso" system sounds on Mac OS). A new trial started either immediately after the response or after 4 second when the observer did not respond.

In each run, observers were shown five different contrast intensities. Each intensity was presented either 50 or 60 times (we increased stimulus repetitions after observing that longer run durations did not impair the performance). Stimulus presentation order was pseudorandomized and each contrast intensity appeared on the left or right side equal number of times. Each run lasted around 10 min, and observers took short breaks between the runs by disengaging their head from the chin rest while remaining in the room when possible. Subjects completed different number of runs in different conditions. In the condition with random trial order, 15 subjects at RIKEN completed an average of 9 runs (range, 5–13; mean number of trials per subject, 2,510), whereas 9 subjects at Stanford completed an average of 3 runs (or 900 trials). At RIKEN, 8 of the observers from the initial group of 15 also took part in subsequent bias induction experiments. They completed an average of 11 runs across all bias induction conditions with an average number of trials being 3,298. At Stanford, all nine subjects also participated in bias induction experiments. Each subject finished 12 runs across all bias induction conditions comprising an average of 3,600 trials. Each bias induction condition was typically completed on different testing days, and the order of conditions was randomized and counterbalanced across subjects.

To test whether we could induce choice history biases we manipulated trial order such that 80% of time after a success or failure the stimulus side stayed or switched sides. To do this, we initially generated a random sequence of trials. Stimulus side on a trial was then chosen to stay or switch 80% of the time to match the desired trial order statistic given the choice outcome of the previous trial. If the stimulus side did not match the random pregenerated order, a future matching trial was swapped with the current trial. Because stimuli were pregenerated in a random order with equal number of stimulus presentations on either side and contrast intensities, the stimulus swapping guaranteed that subjects were presented stimuli on the same side an equal number of time. Also, this ensured that the run size in condition with random trial order did not differ from conditions in which the trial order statistics were manipulated. Stimulus swapping became impossible for a few trials at the end of some runs, but this resulted in negligible changes of desired trial sequence probabilities.

To learn the task and response buttons, subjects were given a few practice trials. During practice, subjects also learnt that the task difficulty varies with stimulus contrast and were instructed to guess when in doubt. The same set of instructions were given in all conditions. We did not discuss the structure of trials, such as randomization or trial dependencies.

UCL. Subjects were asked to detect vertical gratings (1.7 c/°, 4.8° in width and height) presented randomly either to the left or right of a white fixation dot for 200 ms. The grating was embedded in a white noise patch, and noise patches were presented both on the right and left of fixation (Fig. 1*B*). After the stimulus offset, a high contrast white noise pattern appeared in place of the stimulus as a mask to prevent image aftereffects. The monitor was situated 120 cm from the subject and had a uniformly gray background which was set to the midpoint of the calibrated monitor's full range. Stimuli were presented using the method of constant stimuli wherein grating contrasts were randomly selected from values of 0% (no stimulus), 1%, 2%, 3%, 4%, or 8%.

Subjects were asked to fixate and then detect where the stimulus was presented. They were warned about the task difficulty and encouraged to go with their instinct to guess the stimulus side when uncertain. After the grating was masked by high contrast noise, subjects could report the detected stimulus by pressing the Z or M keys for left or right side, respectively. The response was followed by a brief visual feedback. The fixation dot changed to black when subjects made a mistake, but otherwise remained white. The next trial started after the feedback allowing subjects to set their own pace of the experiment. Each run consisted of 550 trials (50 trials per contrast), and each subject completed three runs.

Analysis.

Data processing. Unanswered trials were excluded. In the natural history bias condition, 51 such trials were excluded across all subjects, which was 0.13% of the total number of trials in that condition. In all induced bias conditions, there were 13 unanswered trials which was 0.04% of the total number of trials in those conditions.

Psychophysical analysis. For each subject and for each run, proportion of rightward choices for each contrast intensity were computed as

$$p_R(c)=\frac{N_R(c)}{N_R(c)+N_L(c)}, \qquad [1]$$

where c is the contrast intensity, and $N_R(c)$ and $N_L(c)$ are the total number of rightward and leftward choices for contrast c, respectively. To get the mean proportion of rightward responses for each subject, we averaged proportions computed for each run.

Probabilistic choice model. To quantify the influence of stimulus contrast and previous trial outcome (success or failure) on current trial choices, we used a probabilistic choice model (14). The model is a binomial logistic regression which estimates the probability of selecting the right or left side based on weighting of the stimulus location (separately for each contrast), success and failure outcome on the previous trial and overall bias (preference of one side over the other). The probabilistic choice model assumes that the log-odds of a probability of choosing a stimulus on the right (p) or left ($1-p$) is a linear function of sensory (stimulus contrast) and nonsensory (choice history) parameters

$$L(t)=\ln\left[\frac{p_R(t)}{p_L(t)}\right]=\ln\left[\frac{p_R(t)}{1-p_R(t)}\right]=\beta_{Success}C_{Success}(t-1)+\beta_{Fail}C_{Fail}(t-1)+\beta_1 I_1(t)+\ldots+\beta_n I_n(t)+\beta_{L/R}, \qquad [2]$$

where t is the trial number, P_R and P_L are probabilities of choosing right and left, and $C_{Success}$ and C_{Fail} are the subject choice on the previous trial contingent on whether it was a success or failure encoded as −1 and 1 for left and right choices, respectively. For previous trials that were not a success or failure, both $C_{Success}$ and C_{Fail} were set to 0, and this only applied to the first trial of each run (unanswered trials, which were subject to the same rule, were excluded). I is the stimulus intensity encoded as −1 and 1 to indicate left and right stimulus side and 0 when the corresponding stimulus intensity was not presented. $\beta_{Success}$, β_{Fail}, $\beta_1\ldots\beta_n$ are weights that were obtained by fitting the model, and $\beta_{L/R}$ is the weight of the model intercept that indicates a general left or right bias. Eq. **2** is a continuous function that ranges from $-\infty$ to $+\infty$. To convert log-likelihood estimate into a probability, the logit (in generalized form called softmax) function was used

$$p_R(t)=\frac{1}{1+e^{-L(t)}}, \qquad [3]$$

where p_R is the probability of choosing right rising from 0 (choose left) to 1 (choose right), and $L(t)$ is the log-likelihood of choosing right computed using Eq. **2**.

Recently, Frund et al. (18) presented a model that captures choice history biases by representing choices as a combination of stimulus (i.e., left/right presentation) and response (e.g., left/right response) weights. This parametrization of choice history biases allows distinguishing between response biases, such as when subjects generally prefer to switch independent of feedback as well as stimulus driven choice history biases, such as when subject's choices are driven by the stimulus on the previous trial. Combination of these parameters can help identify win-stay/lose-switch biases. The Frund et al. (18) model has the same number of parameters as our model and is therefore equivalent. Indeed, we found no difference in the fitted log-likelihoods of the two models. Preference for using one or the other models would depend on which parameterization is more conveniently interpreted in the context of the questions being asked.

Model Fitting. We fitted the probabilistic choice model to data from each run. For model fitting, we built a matrix in which the columns consisted of sensory parameters, nonsensory parameters, and subject choice histories encoded as −1 or 1 to indicate left or right and 0 to code unrelated parameters in the current trial. The number of sensory columns matched the number of stimulus contrast intensities used during the run (typically 5). Nonsensory parameters consisted of two columns indicating success or failure on the previous trial. For each trial, one of these two columns was set to either −1 or 1, whereas the other column was 0. Values of −1 or 1 indicated the choice on the previous trial (left or right) that led to failure or success. Finally, subject choices were coded in a separate column and values of −1 or 1 indicated left or right choices, respectively.

To prevent overfitting, which could happen in runs in which subjects responded 100% of the time to high-contrast stimuli, we used "ridge" regularized logistic regression (58). Ridge (or L2) logistic regression has an error parameter that is modeled to include the squared sum of regression coefficients scaled by an additional parameter Λ

$$L_{regularized}=L+\Lambda\sum\beta^2, \qquad [4]$$

where L is defined by Eq. **2**, β are parameters defined in Eq. **2** with the exception of $\beta_{L/R}$, and $\Lambda\geq 0$ is a free parameter that controls "shrinkage" of β parameter estimates. To estimate an optimal Λ, we used a cross-validation procedure. For each run, we fitted the model to randomly sampled 80% of the trials and validated the model on 20% of the remaining trials (59). This random sampling and validation procedure was performed 100 times using 19 values of Λ, which were exponentially distributed from e^{-8} to 20 (thus, for each Λ, 100 values of likelihood were computed). We then selected the Λ value that produced the highest likelihood. The likelihood was computed as

$$LH=\prod z(t), \qquad [5]$$

where t refers to trial number and $z(t)$ corresponds to

$$z(t)=\left\{\begin{array}{ll}\lambda+(1-2\lambda)p_R(t), & \textit{if subject choice}=\textit{right}\\ \lambda+(1-2\lambda)[1-p_R(t)]), & \textit{if subject choice}=\textit{left}\end{array}\right\},$$

where $p_R(t)$ was the probability of choosing right computed using Eq. **3**. λ denotes the lapse rate, which is the proportion of mistakes when the stimulus is clearly visible, such as missing the stimulus due to an eye blink. The lapse rate was computed as a proportion of errors at the highest stimulus intensity as proposed by Prins (60). This approach to lapse rate computation was possible because we chose the highest contrast such that it was clearly visible. The model was fitted in R (61) using the "glmnet" package (62). We note that the model is generalizable and can be fitted using the probit function instead of logit (Eq. **3**). Because logit and probit functions are nearly identical, we expect the difference between logistic and probit regressions to be indistinguishable.

In addition to L2 regularization, we also fitted the model using L1 regularization. The difference between the two approaches is that, although L2 tends to affect both strong and weak weights, L1 regularization tends to redistribute the bias toward strong weights (e.g., weights corresponding to strong contrast intensity) and eliminate model parameters that are weak. We found that both L2 and L1 regularizations produce similar results, which testifies to the robustness of choice history biases.

Interpreting Model Weights. Positive weights ($\beta_{Success}$ and β_{Fail}) obtained for success and failure choice history biases indicated that subjects preferred to remain on the same side they chose on the previous trial, whereas negative weights indicated a preference to switch sides. Positive weights obtained for contrasts ($\beta_1 ... \beta_n$) indicated that subjects, overall, correctly chose the side where stimuli of a given contrast were presented, whereas negative weights indicated that subjects chose the side opposite to where those stimuli were presented (making mistakes). Finally, positive and negative values of the intercept ($\beta_{L/R}$) indicated left or right choice biases, respectively.

Model Selection. We used likelihood ratio tests, which allows comparisons between two nested models (63), to evaluate whether subjects had choice history biases. We compared the full model (probabilistic choice model) to a simpler no-history model that contained the same parameters but without choice history weights. Parameters of the full model were stimulus contrast weights, choice history weights of success and failure on the previous trial, and a weight for general bias for left or right choices. For this procedure, we first computed the likelihood ratio statistic LR as follows:

$$LR = -2\log\left(\frac{LH_{Full\ Model}}{LH_{No-history\ Model}}\right),$$

where $LH_{Full\ Model}$ and $LH_{No-history\ Model}$ were maximum likelihoods of the full and no-history models, respectively, computed using Eq. **5**. The distribution of values of the likelihood ratio statistics asymptotes to a χ^2 distribution with degrees of freedom equal to the difference in the number of parameters between the two models, which in our case was 2 (success and failure bias parameters were removed from the no-history model). Second, we subjected the likelihood ratio to χ^2 statistics to compute the probability of the null hypothesis that the no-history model was better. P values less than or equal to 0.05 were used to reject the null hypothesis and accept the full model as the better fitting model.

Estimating Sensitivity Loss. To estimate differences in subjects' performance with and without biases, we fitted psychometric curves to the model simulated choices with and without bias. In one case, the model simulated subjects' biased choices by including both sensory and nonsensory terms. In another case, the model included only sensory terms and all nonsensory terms were set to zero to simulate a subject who has no biases and only responds to sensory signals. In a simulated run, both biased and unbiased models made choices on the same set of randomly generated trials of different contrast intensities presented to the left or right visual field. Each contrast intensity was presented 48 times, similar to the number of trials used during the experiment. We next fitted psychometric function to the proportion of rightward choices using Eq. **6** to extract the slope as a measure of sensitivity. The sensitivity loss was computed as the ratio of two slopes that indicated the percent of drop of slope steepness caused by choice history biases. Each run was simulated 500 times for each subject, and the median decline of sensitivity is shown in Fig. 4*A*. Each subject model was constructed by taking the average weights of multiple runs collected during the experiment.

Fitting Psychometric Curves. The psychometric function models the relationship between the stimulus intensity and subject's responses. For each run, the fit was applied to the proportion of rightward choices when stimuli were presented to the right or left (Fig. 2). This approach to fitting psychometric curves is more general as it helps to estimate left/right choice biases (64). The psychometric function had the following form:

$$\Psi(x; \alpha, \beta, \lambda) = \lambda + (1 - 2\lambda)F(x; \alpha, \beta), \quad [6]$$

where α and β are free parameters that correspond to the threshold and slope of the psychometric function. λ is the lapse rate that indicates unintentional errors that occur when the stimulus is obvious. The lapse rate was computed as a proportion of errors when the stimulus was clearly visible (60). The stimulus was clearly visible at 3% contrast for data collected at RIKEN and Stanford and 8% for the data acquired at the UCL (mean $\lambda = 0.02$; 95% CI, 0.01–0.03). The same lapse rates were also used in estimating the weights of the probabilistic choice model. Low lapse rates confirm that subjects were alert to the task and choice history biases were not a result of inattention but rather unintended misses. We applied probit analysis in which the psychometric function, $F(x; \alpha, \beta)$, was a cumulative Gaussian

$$F(x; \alpha, \beta) = \frac{1}{\sqrt{2\pi}\beta} \int_{-\infty}^{contrast} e^{-\frac{(x-\alpha)^2}{2\beta^2}}, \quad [7]$$

We chose probit analysis because it provided a better fit compared with the logistic function (65). To assess changes in subjects' performance, we computed the slope and 75% contrast threshold of psychometric functions. The threshold was computed as a contrast increment that allowed subjects to reach 75% performance from 50% performance.

Variance Inflation Factor. Changing trial order statistics introduces relationship between previous trial and current trial and potentially between choice history bias parameters and stimulus contrast parameters, such that our predictor variables could be determined by other predictors in the model. This collinearity could affect accuracy of estimating model weights (66), which is computed as $1/(1-R^2)$, where R^2 correlation is computed by regressing each parameter of the model with the remaining parameters (67). When R^2 is close to 0 (little correlation between parameters), variance inflation factor (VIF) is close to 1, which indicates no collinearity effects, whereas larger R^2 values lead to larger VIF, and values above 5 mark a collinearity problem.

To ensure that we were actually measuring induced biases and that our trial order manipulations did not spuriously cause changes in fitted choice history weights, we ran several control analyses. We evaluated whether changing trial order statistics introduced collinearity among model weights by computing VIF and found no collinearity effects in our data ($M_{vif} = 1.07$; 95% CI, 1.06–1.08). We also permuted (scrambled) subjects' responses within each run 10 times and fitted the model to these data to estimate choice history biases. This approach preserved trial sequence structure generated through trial order manipulations, but eliminated subjects' choice history biases. We found that choice history weights, as expected, were not different from 0 (mean $\beta_{Fail} = -0.003$, $P = 0.1$; mean $\beta_{Success} = -0.002$, $P = 0.2$). We ran another validation by simulating an observer based on subjects' model weights after removing choice history biases (history weights were set to 0) and found that our trial order manipulations did not induce any artificial biases in the induced direction (all $P > 0.05$, one-tailed t test comparison with 0).

ACKNOWLEDGMENTS. We thank Sivaramakrishnan R. Kaveri, Steeve Laquitaine, and Tancy (Chien-Hui) Kao for valuable discussions. This research was funded by Grants-in-Aid for Scientific Research 24300146, Japanese Ministry of Education, Culture, Sports, Science and Technology. M.C. holds the GlaxoSmithKline/Fight for Sight Chair in Visual Neuroscience.

1. Nowak M, Sigmund K (1993) A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* 364(6432):56–58.
2. Sutton SR, Barto AG (1998) *Reinforcement Learning* (MIT Press, Cambridge, MA).
3. Daw ND, Doya K (2006) The computational neurobiology of learning and reward. *Curr Opin Neurobiol* 16(2):199–204.
4. Lau B, Glimcher PW (2005) Dynamic response-by-response models of matching behavior in rhesus monkeys. *J Exp Anal Behav* 84(3):555–579.
5. Wang Z, Xu B, Zhou H-J (2014) Social cycling and conditional responses in the Rock-Paper-Scissors game. *Sci Rep* 4:5830.
6. Barraclough DJ, Conroy ML, Lee D (2004) Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci* 7(4):404–410.
7. Nash JF (1950) Equilibrium points in N-person games. *Proc Natl Acad Sci USA* 36(1): 48–49.
8. Worthy DA, Hawthorne MJ, Otto AR (2013) Heterogeneity of strategy use in the Iowa gambling task: A comparison of win-stay/lose-shift and reinforcement learning models. *Psychon Bull Rev* 20(2):364–371.
9. Fernberger SW (1920) Interdependence of judgments within the series for the method of constant stimuli. *J Exp Psychol* 3(2):126–150.
10. Senders VL, Sowards A (1952) Analysis of response sequences in the setting of a psychophysical experiment. *Am J Psychol* 65(3):358–374.
11. Verplanck WS, Collier GH, Cotton JW (1952) Nonindependence of successive responses in measurements of the visual threshold. *J Exp Psychol* 44(4):273–282.
12. McGill WJ (1957) Serial effects in auditory threshold judgments. *J Exp Psychol* 53(5): 297–303.
13. Tune GS (1964) Response preferences: A review of some relevant literature. *Psychol Bull* 61(4):286–302.
14. Busse L, et al. (2011) The detection of visual contrast in the behaving mouse. *J Neurosci* 31(31):11351–11361.
15. Carandini M, Silva L, Busse L, Dakin S (2012) Vision and superstition in mouse and man. *J Vis* 12(9):620.
16. Raviv O, Ahissar M, Loewenstein Y (2012) How recent history affects perception: The normative approach and its heuristic approximation. *PLOS Comput Biol* 8(10):e1002731.
17. Abrahamyan A, Gardner JL (2013) Past failures bias human decisions (Salt Lake City). Available at cosyne.org/cosyne13/Cosyne2013_program_book.pdf. Accessed May 21, 2016.
18. Fründ I, Wichmann FA, Macke JH (2014) Quantifying the effect of intertrial dependence on perceptual decisions. *J Vis* 14(7):9–9.

19. Yeshurun Y, Carrasco M, Maloney LT (2008) Bias and sensitivity in two-interval forced choice procedures: Tests of the difference model. *Vision Res* 48(17):1837–1851.
20. Raviv O, Lieder I, Loewenstein Y, Ahissar M (2014) Contradictory behavioral biases result from the influence of past stimuli on perception. *PLOS Comput Biol* 10(12):e1003948.
21. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113(4):700–765.
22. Blackwell HR (1952) Studies of psychophysical methods for measuring visual thresholds. *J Opt Soc Am* 42(9):606–616.
23. Rabin M, Vayanos D (2010) The gambler's and hot-hand fallacies: Theory and applications. *Rev Econ Stud* 77(2):730–778.
24. Miller JB, Sanjurjo A (2015) Surprised by the gambler's and hot hand fallacies? A truth in the law of small numbers. *IGIER Working Paper #552*. Available at ssrn.com/abstract=2627354. Accessed May 21, 2016.
25. Verplanck WS, Cotton JW, Collier GH (1953) Previous training as a determinant of response dependency at the threshold. *J Exp Psychol* 46(1):10–14.
26. Fechner GT (1966) *Elements of Psychophysics* (Holt, Rinehart and Winston, New York).
27. Nachmias J, Sansbury RV (1974) Letter: Grating contrast: Discrimination may be better than detection. *Vision Res* 14(10):1039–1042.
28. Legge GE (1981) A power law for contrast discrimination. *Vision Res* 21(4):457–467.
29. Tervo DGR, et al. (2014) Behavioral variability through stochastic choice and its gating by anterior cingulate cortex. *Cell* 159(1):21–32.
30. Pelli DG, Bex P (2013) Measuring contrast sensitivity. *Vision Res* 90:10–14.
31. Worthy DA, Otto AR, Doll BB, Byrne KA, Maddox WT (2015) Older adults are highly responsive to recent events during decision-making. *Decisions* 2(1):27–38.
32. Körding K (2007) Decision theory: What "should" the nervous system do? *Science* 318(5850):606–610.
33. May KA, Solomon JA (2015) Connecting psychophysical performance to neuronal response properties II: Contrast decoding and detection. *J Vis* 15(6):9–9.
34. Chirimuuta M, Tolhurst DJ (2005) Does a Bayesian model of V1 contrast coding offer a neurophysiological account of human contrast discrimination? *Vision Res* 45(23):2943–2959.
35. Goris RLT, Putzeys T, Wagemans J, Wichmann FA (2013) A neural population model for visual pattern detection. *Psychol Rev* 120(3):472–496.
36. Pelli DG (1985) Uncertainty explains many aspects of visual contrast detection and discrimination. *J Opt Soc Am A* 2(9):1508–1532.
37. Vintch B, Gardner JL (2014) Cortical correlates of human motion perception biases. *J Neurosci* 34(7):2592–2604.
38. Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. *Nat Neurosci* 5(6):598–604.
39. Girshick AR, Landy MS, Simoncelli EP (2011) Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nat Neurosci* 14(7):926–932.
40. Schwartz O, Hsu A, Dayan P (2007) Space and time in visual context. *Nat Rev Neurosci* 8(7):522–535.
41. Fischer J, Whitney D (2014) Serial dependence in visual perception. *Nat Neurosci* 17(5):738–743.
42. Guo K, et al. (2004) Effects on orientation perception of manipulating the spatio-temporal prior probability of stimuli. *Vision Res* 44(20):2349–2358.
43. Sotiropoulos G, Seitz AR, Seriès P (2011) Changing expectations about speed alters perceived motion direction. *Curr Biol* 21(21):R883–R884.
44. Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J Neurosci* 31(17):6339–6352.
45. Mulder MJ, Wagenmakers EJ, Ratcliff R, Boekel W, Forstmann BU (2012) Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *J Neurosci* 32(7):2335–2343.
46. Rao V, DeAngelis GC, Snyder LH (2012) Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. *J Neurosci* 32(29):10063–10074.
47. Sagi D (2011) Perceptual learning in vision research. *Vision Res* 51(13):1552–1566.
48. Donahue CH, Seo H, Lee D (2013) Cortical signals for rewarded actions and strategic exploration. *Neuron* 80(1):223–234.
49. Ito M, Doya K (2009) Validation of decision-making models and analysis of decision variables in the rat basal ganglia. *J Neurosci* 29(31):9861–9874.
50. Ting CC, Yu CC, Maloney LT, Wu SW (2015) Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference. *J Neurosci* 35(4):1792–1805.
51. Summerfield C, Tsetsos K (2012) Building bridges between perceptual and economic decision-making: Neural and computational mechanisms. *Front Neurosci* 6:70.
52. Foster KR, Kokko H (2009) The evolution of superstitious and superstition-like behaviour. *Proc Biol Sci* 276(1654):31–37.
53. Killeen PR (2008) Learning as causal inference. *Quantitative Analyses of Behavior*, eds Common ML, Nevin JA (Ballinger Publishing, Cambridge, MA), pp 92–112.
54. Tversky A, Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.
55. Lee D (2006) Neural basis of quasi-rational decision making. *Curr Opin Neurobiol* 16(2):191–198.
56. Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10(4):433–436.
57. Pelli DG (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spat Vis* 10(4):437–442.
58. Hastie T, Tibshirani R, Friedman J (2009) *The Elements of Statistical Learning* (Springer Science & Business Media, New York).
59. Conway D, White JM, White JM (2012) *Machine Learning for Hackers* (O'Reilly Media, Inc., Sebastopol, CA).
60. Prins N (2012) The psychometric function: The lapse rate revisited. *J Vis* 12(6):25.
61. R Core Team (2014) R: A language and environment for statistical computing. Available at www.R-project.org. Accessed May 21, 2016.
62. Friedman J, Hastie T, Tibshirani R (2010) Regularization paths for generalized linear models via coordinate descent. *J Stat Softw* 33(1):1–22.
63. Casella G, Berger RL (2002) *Statistical Inference* (Duxbury, Pacific Grove, CA), 2nd Ed.
64. Klein SA (2001) Measuring, estimating, and understanding the psychometric function: a commentary. *Percept Psychophys* 63(8):1421–1455.
65. Wichmann FA, Hill NJ (2001) The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept Psychophys* 63(8):1293–1313.
66. Kleinbaum DG, Klein M, Pryor ER (2002) *Logistic Regression: A Self-Learning Text* (Springer, New York).
67. Fox J, Weisberg S (2011) *An R Companion to Applied Regression* (SAGE, Thousand Oaks, CA).