The challenge of measuring long-term positive aftereffects

Gerrit W. Maus[1], Wesley Chaney[2], Alina Liberman[3], and David Whitney[1,2,3]

Adaptation is one of the longest-studied phenomena in perception and neuroscience. Adaptation generally results in negative perceptual aftereffects: after prolonged exposure to a specific feature, perception of a neutral stimulus is biased in the opposite direction [1,2]. A recent paper in *Current Biology* [3] challenged this view by proposing that, additionally, adaptation biases perception in the *same* direction as features observed over a relatively long time from the past. This finding challenges dominant theories of visual adaptation; however, here we argue that long-term positive correlations are not due to neural or perceptual processes but arise due to short-term negative aftereffects. Thus, existing models of adaptation remain unchallenged, and critical evaluations of how adaptation could predictively aid perception are still needed.

Chopin and Mamassian [3] presented observers with binocularly rivalrous oriented gratings within series of non-rivalrous gratings [4]. Their analysis correlates observers' responses with stimuli presented in windows of different durations and at different time points (lags) in the past (their Figure 2). Perception of rivalrous gratings was biased opposite to previously shown non-rivalrous gratings. In addition to this negative aftereffect, observers' responses were positively correlated with stimuli from a 'reference window' in the past. The same held true in a tilt-aftereffect experiment [5].

However, short-term negative aftereffects alone account for this pattern of correlations. We simulated an artificial observer whose responses were determined by a noisy short-term negative aftereffect (see Supplemental Information). Performing the same analysis, we found significant positive correlations for large window durations and time lags (Figure 1A), similar to results in human observers (Figure 2 in [3]). Different parameters for aftereffect length and noise yielded similar results (Supplemental Figure S1). Positive correlations arise because of an interaction between the short-term negative aftereffect and random fluctuations in the stimulus sequence. Any random sequence will exhibit fluctuations in the proportion of left or right stimuli. Because of the short-term aftereffect, responses are correlated negatively with the stimuli in recent history and thus show similar fluctuations in counterphase. Increasing the time lag of stimulus windows amounts to shifting the timecourse relative to observers' responses. Because of the fluctuations in both timecourses, some phase shifts will necessarily

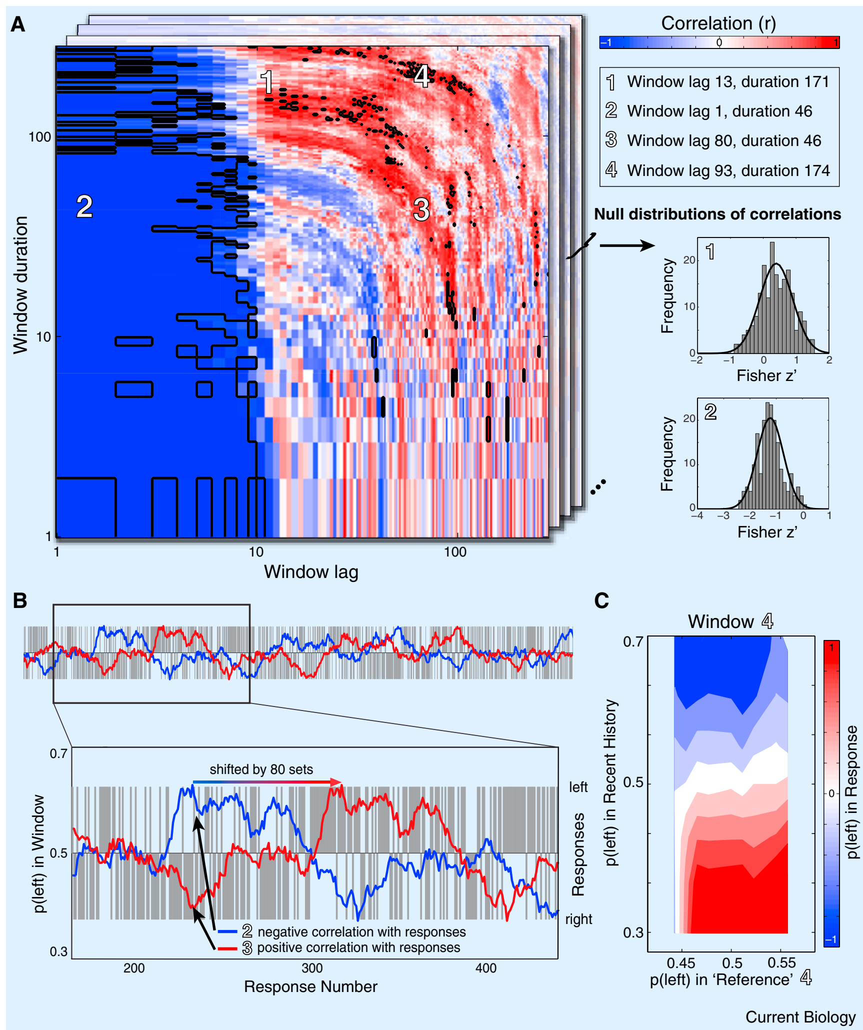

Figure 1. A short-term negative aftereffect can produce seeming capture of responses by long-term stimulus history.
(A) Correlations between responses of simulated observers and stimulus windows of varying durations and lags lead to remarkably similar results as human observers in [3]. Significant correlations are indicated by black outlines. Repeating this simulation many times generates a null distribution for correlations expected from just a negative aftereffect (shown for two example windows on the right; also see Supplemental Figure S2). (B) Simulated timecourses of responses (gray bars) and stimuli in a sliding average window (blue, 2 in A). Because of the short-term aftereffect, responses are negatively correlated with the averaged stimulus timecourse. Increasing the window's lag (red, 3 in A) relative to the responses can turn these negative correlations to positive. (C) Simulated responses as a function of stimuli in the recent history and a selected 'reference window' (4 in A). A positive relationship between reference window and responses (as shown for a selected window in Figure 3 of [3]) can occur by chance due to noise in the observer.

lead to positive correlations (Figure 1B). Indeed, because averaging of long stimulus windows reduces the number of statistically independent samples, mathematical considerations predict more positive than negative correlations (see Supplemental Information).

Demonstrating a genuine long-term positive aftereffect necessitates statistically comparing empirical correlations with a null distribution generated by assuming just a short-term aftereffect (rather than that no correlations exist). We generated such null distributions by repeating our simulation many times on random sequences and found positive correlations for large window durations and lags, all of which resulted from only the short-term negative aftereffect (Figures 1A and Supplemental Figure S2).

To explain the long-term assimilative effect, Chopin and Mamassian [3] proposed a model of how recent stimulus history and a long-term 'reference' window are taken into account in perceptual decision-making. This model predicts effects of both recent history and reference on observers' responses. For a selected reference window they showed such effects (their Figure 3). But the same analysis on simulated data revealed that similar interactions could occur by chance (Figure 1C), even though recent history and reference window do not independently influence responses in our simulations. Again, additional influence of long-term history — beyond that of short-term history — should be assessed by comparison to null distributions from simulations.

More consideration is needed regarding the proposal that long-term positive aftereffects could serve a 'predictive' purpose. Chopin and Mamassian [3] write: "Implicit predictions are based on the assumption that the distributions of orientations should match between recent history and the remote reference" (p. 625). This 'gambler's fallacy' model, however, assumes that the proportion of observable orientations in the world is static and unchanging over the period in question (empirically ~13 minutes for the stimuli in [3]). Considering the dynamic properties of the natural world, one could reasonably argue that the best predictions for the state of the world are based on its current or very recent state, not a remote past reference. Physical auto-correlations, by definition, are strongest at short timescales. To overcome internal perturbations in the perceptual system, there is no reason to believe that an estimate from ~10 minutes ago is any more reliable or less biased than one based on more recent evidence.

Our simulations show that human perception and behavior can exhibit deceptive long-term temporal structure. While negative aftereffects in both rivalry [4] and tilt [5] are well established, the long-term assimilative effects in [3] and our simulation are spurious. Previous models of visual adaptation, including error correction, decorrelation, or Bayesian inference processes [1,2], can easily accommodate the apparent assimilative structure; they need no modification or new parameters.

Supplemental Information

Supplemental Information includes details of experimental procedures and two figures, and can be found with this article online at http://dx.doi.org/10.1016/j.cub.2013.03.024.

Acknowledgments

We thank A. Chopin, R. Denison, J. Fischer, S. Haroz, P. Mamassian, and T. Sweeny for thought-provoking discussions. Supported in part by grants NIH EY007043 (W.C.), NSF 1106400 (A.L.), and NIH EY018216 (D.W.).

References

1. Clifford, C.W.G., Webster, M.A., Stanley, G.B., Stocker, A.A., Kohn, A., Sharpee, T.O., and Schwartz, O. (2007). Visual adaptation: neural, psychological and computational aspects. Vision Res. *47*, 3125–3131.
2. Kohn, A. (2007). Visual adaptation: physiology, mechanisms, and functional benefits. J. Neurophysiol. *97*, 3155–3164.
3. Chopin, A., and Mamassian, P. (2012). Predictive properties of visual adaptation. Curr. Biol. *22*, 622–626.
4. Blake, R., and Overton, R. (1979). The site of binocular rivalry suppression. Perception *8*, 143–152.
5. Gibson, J.J., and Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies. J. Exp. Psychol. *20*, 453–467.

[1]Department of Psychology, [2]Vision Science Program, School of Optometry, [3]Helen Wills Neuroscience Institute, University of California Berkeley, Berkeley, CA 94720, USA.
E-mail: maus@berkeley.edu

CrossMark

Response: Genuine long-term positive aftereffects

Adrien Chopin[1]
and Pascal Mamassian[2]

Perceptual adaptation traditionally leads to negative aftereffects: observers experience the opposite of what they have just been exposed to. In a recent paper in this journal [1], we reported that this negative correlation between the current percept and the recent ones is accompanied by a positive correlation with events occurring further in the past. This result suggests a simple mechanism to recalibrate a sensory system. Events occurring in a remote temporal window can be used to estimate some statistics on the environment, and events occurring recently are then compared to this estimate. A recalibration is initiated when a discrepancy exists between recent and remote statistics. This proposal is very different from the traditional view of adaptation whereby calibration is purely determined by recent events. Maus *et al*. [2] argue that our results can be explained by a simple negative aftereffect model; here, we refute their arguments.

In our recent paper [1], we analysed psychophysical data of binocular rivalry and tilt aftereffect experiments by measuring the correlation between the probability of perceiving an event and the proportion of that event in windows of different sizes and positions in the past. Maus *et al*. [2] suggest that this analysis can lead to positive correlations for remote windows when they simulate an observer who is only subject to the classical negative aftereffect. They propose that positive correlations may arise from a shift of the event proportion time-courses when using a lagged window. Given a negative aftereffect model, responses are negatively correlated with the proportions and, because of the shift, will sometimes be positively correlated. For this to be systematically true, however, fluctuations need to be in counterphase with fluctuations in the other lagged window (as in Figure 1B in [2]), a scenario that would require fluctuations to reverse periodically

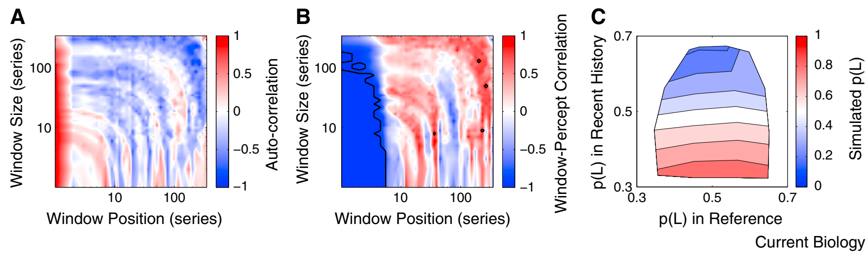

Figure 1. Correlations within stimuli and their evoked percepts.
(A) Auto-correlations within a simulated sequence of stimuli. (B) Correlation between the stimuli in the window and the simulated percept of an aftereffect model for the same sequence of stimuli used to generate (A). Note the inversion of the sign of the correlations relative to (A), except for window positions less than the aftereffect size that are always negative in (B). (C) Probability to obtain percept 'L' in the simulation as a function of the proportions of stimuli 'L' in the recent and remote windows. The data are averaged across all remote windows whose correlation was positive and significant in (B).

with a specific phase matching the lag.

Maus *et al*. [2] also argue that these correlations "arise due to an interaction between the short-term negative aftereffect and random fluctuations in the stimulus sequence". One way to reveal these patterns of correlation within the stimulus sequence is to compute an auto-correlation, that is, to correlate the stimulus at time (t+1) — instead of the response of the observer — with the stimuli within windows of different sizes and positions in the past up to time (t). This auto-correlation analysis does indeed highlight some sporadic correlations in the simulated stimulus sequence (Figure 1A); however, other sequence simulations reveal different correlations. Interestingly, an aftereffect model will closely follow these auto-correlations, up to a sign inversion (Figure 1B). Pearson's correlation between Figure 1A and 1B reveals a strong negative relationship reflecting this sign inversion ($r = -0.54$, $p < 0.001$). The auto-correlations may be more often negative for large windows because of a –1/N bias [3]. These negative auto-correlations would then produce positive correlations with responses in a negative aftereffect model. In contrast, the correlation between our observers' data (shown in Figure 2 in [1]) and the auto-correlations of the stimulus sequence revealed a weak and positive relationship ($r = 0.09$, $p < 0.001$ in the rivalry experiment, and $r = 0.08$, $p < 0.01$ in the tilt aftereffect experiment). In other words, the sporadic auto-correlations inherent to the sequence presentation are revealed by the aftereffect model but are not responsible for the correlation structure in our analysis in [1].

In our original paper [1], we presented a fine analysis of the relation between a percept at time (t) and the stimuli in each window. Our model predicts a diagonal gradient when the probability to obtain the next percept is expressed as a function of the proportion of events in the recent and remote windows. For both experiments we ran, we found evidence for such diagonal patterns (Figure 3B,C in [1]). Maus *et al*. [2] report a similar pattern between the recent window and a selected remote window (Figure 1C in [2]). We performed this analysis for all the significant positive windows generated by the simulated aftereffect model and never found that pattern (see Figure S1 in the Supplemental Information for individual plots and Figure 1C for their average).

In an effort to better assess the influence of remote and recent windows, we perform here a logistic regression with the following equation:

$$\text{Logit}(y) = a + \text{beta}_1.x_1 + \text{beta}_2.x_2$$

with y the binary percept, x_1 and x_2 the proportions of 'Left' events in recent and remote windows. For our original experiments [1], the regression led to significant effects of the proportion of events in the recent and remote windows (Supplemental Table S1). For data simulated from the aftereffect model, the logistic regression revealed, as expected, a significant influence of the recent window on the simulated percept, but importantly no influence of the remote windows that were significant and positive in the correlation analysis (Supplemental Table S1).

Lastly, Maus *et al*. [2] criticize our model's ecological validity, because it assumes that "the proportion of observable orientations in the world is static and unchanging over [a period of 13 minutes]". Our model only assumes that, in the particular setting of our experiments, the number of samples to reliably estimate orientation statistics was 300 and the number of samples to reliably estimate the current distribution of orientations was 100. Interestingly, our model also allows a self-calibration to overcome internal perturbations of the sensors (error correction [4]). A model that relies purely on recent events and that does not compare the distribution of these events with a norm is unable to reach this goal.

In summary, we do agree that there are sporadic auto-correlations in our stimulus sequences; however, these auto-correlations do not explain the correlations we found with the observers' responses. The aftereffect model cannot account for the diagonal pattern found in the responses when expressed as a function of recent and remote windows. The logistic regression confirmed that the remote window proportions can account for additional variance in the responses that is not explained by the negative aftereffect model.

Supplemental Information

Supplemental Information includes experimental procedures, one figure and one table and can be found with this article online at http://dx.doi.org/10.1016/j.cub.2013.02.025

Acknowledgments

We thank Simon Barthelmé for constructive comments.

References

1. Chopin, A., and Mamassian, P. (2012). Predictive properties of visual adaptation. Curr. Biol. *22*, 622–626.
2. Maus, G.W., Chaney, W., Liberman, A., and Whitney, D. (2013). The challenge of measuring long-term positive aftereffects. Curr. Biol. *23*, R438–R439.
3. Chatfield, C. (1989). The Analysis of Time Series: An Introduction (4th Edition). London: Chapman & Hall.
4. Andrews, D.P. (1964). Error-correcting perceptual mechanisms. Q. J. Exp. Psychol. *16*, 104–115.

[1]Laboratoire Psychologie de la Perception, Université Paris Descartes, Paris, France, and Department of Experimental Psychology, University of Oxford, Oxford, UK.
[2]Laboratoire Psychologie de la Perception, Université Paris Descartes, Paris, France, and CNRS, UMR 8158, Paris, France.
E-mail: adrien.chopin@gmail.com, pascal.mamassian@parisdescartes.fr

Supplemental Information: The challenge of measuring long-term positive aftereffects

Gerrit W. Maus, Wesley Chaney, Alina Liberman, and David Whitney

Supplemental Experimental Procedures

Details of the simulated stimuli

We generated a pseudo-random sequence of ‘stimuli’ consisting of -1 and 1. Stimuli were grouped into sets of 1-4 stimuli, using only 22 selected combinations of -1 and 1 as in the study by Chopin and Mamassian (see their Figure S1). In their study, each participant viewed each of these sets 32 times. To produce an equivalent amount of data as in their study with 8 participants, our stimulus sequence consisted of 32 × 8 repetitions of each set, for a total of 16,896 stimuli. To simulate runs with single observers and produce a null distribution of results expected from only a short-term aftereffect (Figure S2), we used shorter sequences with 32 repetitions of each set with 2,112 stimuli in total (equal in length to the sequences shown to single observers by Chopin and Mamassian).

Details of the artificial observer

We implemented an artificial observer that produced one response for each stimulus set. The observer integrated stimuli by taking the mean of all presented stimuli during a period leading up to the current response. The length of this period, the *aftereffect length*, was varied for different observer models between 12 and 100 stimuli (see Figure S1). We also simulated an observer with a ‘leaky integrator’ type aftereffect, where stimuli were linearly weighted by their position in the stimulus sequence, with more recent stimuli weighted more heavily (Figure S1, bottom row). Noise was added to the integration stage by generating random values between -1 and 1, and computing a weighted average between noise and stimulus mean (weighted by the *noise weight*). The observer’s response was 1, if the weighted mean of integrated stimuli and noise was < 0, and -1 if it was > 0, generating a negative aftereffect. For Figure 1 in the main text, we used an aftereffect length of 36 stimuli and noise weighted by 0.33. A range of parameters for aftereffect length and noise weight produced qualitatively the same results as shown in Figure 1A (see Figure S1).

Details of the analysis

We performed the same analysis as Chopin and Mamassian on our simulated stimuli and the artificial observer’s responses. For each window of a certain lag and duration, stimulus proportions were sorted into 9 bins and correlated with the proportions of corresponding responses. Confidence intervals were calculated by a bootstrap procedure—resampling the original bins 5000 times with replacement. Significant correlations that survived a Bonferroni correction for the number of windows are marked by black outlines in Figure 1A.

It should be noted that because of overlapping stimulus windows, the assumption of independent samples for Pearson's correlation is violated, leading to inflated correlation coefficients. The bootstrap procedure described above resamples binned values from the same overlapping windows, and thus does not correct for this problem. A more appropriate statistical test would be to compare human observer data to a null distribution generated by repeatedly shuffling the stimulus sequence (i.e., constructing different random sequences with the same counterbalancing procedure).

To construct a null distribution of correlations expected from just a short-term negative aftereffect, we calculated the pattern of correlations for 200 different pseudo-random sequences and simulated responses (equal in length to those presented to single human observers by Chopin and Mamassian). Different sequences caused different patterns of positive and negative correlations in the windows, yet on average the analysis still produces positive correlations for large windows durations and lags (see next section on why there are more positive than negative correlations). Figure S2A shows 9 individual iterations, the mean correlations for 200 different sequences are shown in Figure S2B. Mean correlations are calculated on Fisher z' transformed correlation coefficients and then transformed back to Pearson's r. The distribution of correlation coefficients for a few exemplar window lags and durations is shown in Figure 1A and S2C. To test the hypothesis that long-term stimulus history has an effect on human observers' responses over and above a short-term negative aftereffect, empirical correlation values from human observers should be compared to this sort of null distribution.

How do positive correlations arise?

Figure 1B in the main text provides an intuition of why positive correlations occur in the analysis of stimulus windows and responses. Please note that this Figure is only intended to provide an intuition about how positive correlations can arise, not to suggest that this is how the data were actually analyzed. Averaging stimuli within a sliding window of a given duration amounts to low-pass filtering the stimulus timecourse. Although the stimulus sequences employed by Chopin and Mamassian were balanced for their local history, low-pass filtering results in slow fluctuations of left and right biases in the timecourse. Matching the size of the averaging window to the "size" of the negative aftereffect will result in an almost perfect negative correlation; stimulus and response timecourses effecttively fluctuate in counterphase. In our simulation we added a noise term to the observer model, so correlations were never perfect (i.e., correlations never reached values of -1).

Increasing the time lag of the averaging window is equivalent to temporally shifting the stimulus timecourse relative to the response timecourse. Since the original timecourses are fluctuating in counterphase (as described above), shifting one relative to the other will result in positive correlations for some phase lags. Within the range of time lags analyzed by Chopin and Mamassian (up to ~300), we found both positive and negative correlations, leading to alternating "stripes" of red and blue in the plots of correlations for increasing window lags (e.g. Figure S2A), as expected for autocorrelations of periodically fluctuating signals. However, if the simulation is repeated hundreds of times, on average there are more positive than negative correlations for large windows and long lags (Figure S2B).

Why are there (on average) more positive than negative correlations? The zero-lag correlation is determined to be highly negative because of the negative short-term aftereffect (as described above). Because the two signals are not perfectly periodical, all other negative correlations are expected to be smaller than the zero-lag correlation, whereas no such restriction applies for positive correlations. An additional possible explanation for this empirical finding in our simulation is as follows[1]: Autocorrelations for discrete timeseries of N independent observations are expected to be negatively biased for phase lags unequal to zero, with an expected autocorrelation of 1/N (e.g. chapter 4 in [S1]). For large values of N, this bias should be negligible. However, low-pass filtering of the stimulus timecourse (taking the average in a sliding window) reduces the number of independent observations considerably, especially for large window sizes. Hence, autocorrelations of the stimulus timecourse averaged over a large stimulus window are expected to be negatively biased, and therefore correlations with responses (based on a negative aftereffect) are expected to be positively biased.

For these considerations, the direction of the window lag is actually irrelevant: current responses are expected to be positively correlated with windows in the remote past and *future*. This, of course, cannot possibly be an effect of perceptual adaptation, but had Chopin and Mamassian analyzed both past and future stimuli's correlation with responses, it should have led them to the conclusion that both past and future windows serve as a reference for perception.

While this serves to show the possible origin of artifactual positive correlations between long-term stimulus history and observer responses, the important point is that a simple simulation can be used to generate a null distribution of correlations expected from just a short-term aftereffect, and empirical data should be compared to such distributions.

Supplemental References

S1. Chatfield, C. (1989) The analysis of time series: An introduction (4th Edition). London: Chapman & Hall.

[1] We thank one reviewer for this important insight.

Figure S1

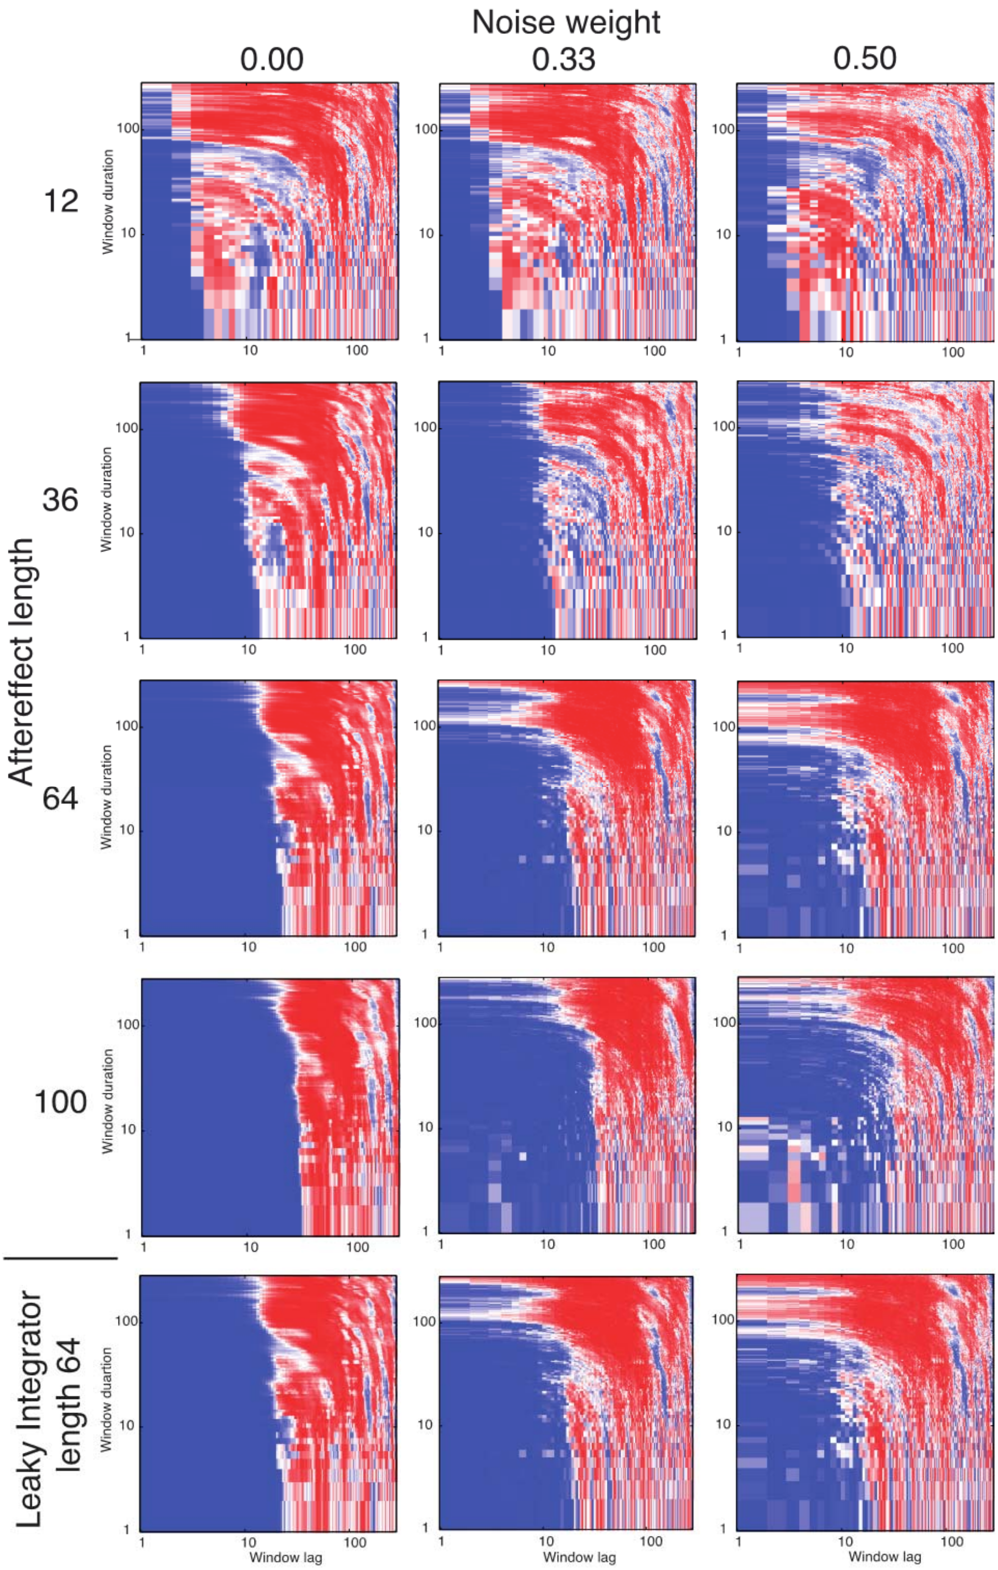

Noise weight
0.00
0.33
0.50
Aftereffect length
12
36
64
100
Leaky Integrator length 64
Window duration
Window duartion
Window lag
1
10
100

Figure S1

Simulation results for a range of parameters of the artificial observer model, all based on the same pseudo-random stimulus sequence.

The observed pattern of correlations occurs consistently for a wide range of artificial observers with varying aftereffect lengths and noise weights in the integration process. Generally, longer aftereffects cause more negative correlations for windows with increasingly longer lags. Regardless of the aftereffect length, however, positive correlations occur for long window durations and long lags. The bottom row shows the same analysis for an artificial observer with a leaky integrator type aftereffect (see Supplemental Experimental Procedures).

Figure S2

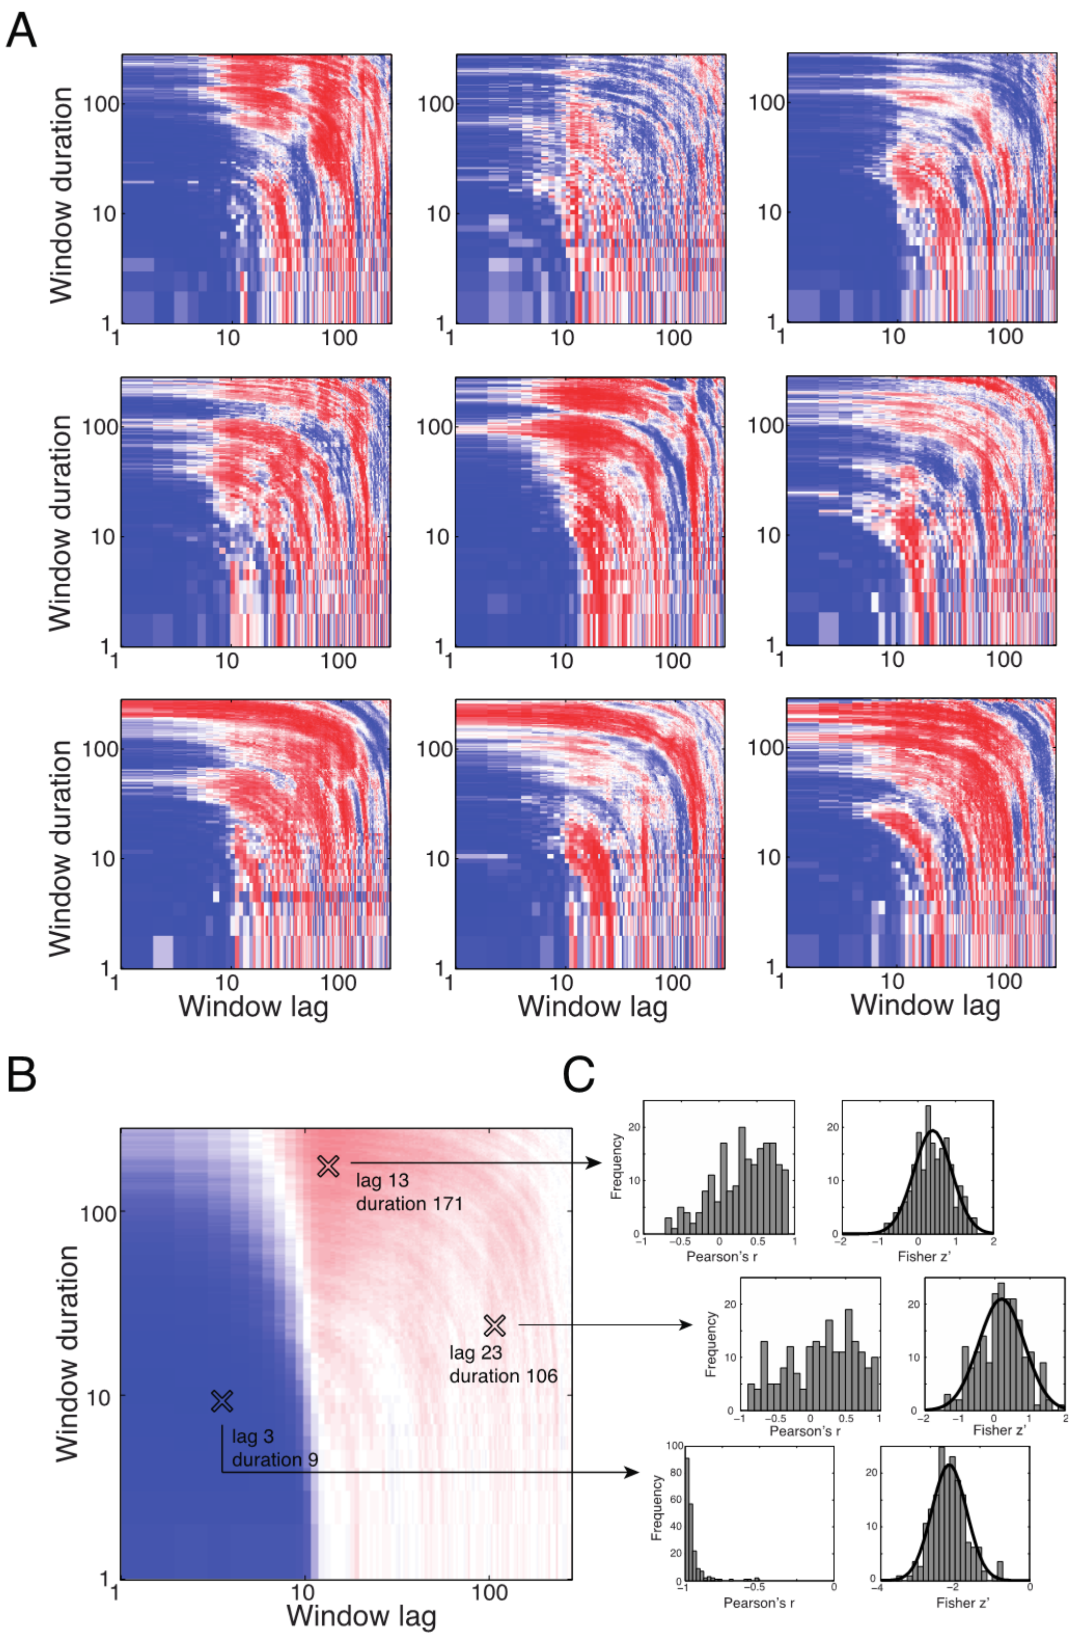

A
Window duration
Window lag
100
10
1
B
lag 13
duration 171
lag 23
duration 106
lag 3
duration 9
C
Frequency
Pearson's r
Fisher z'

Figure S2

Simulation results for different stimulus sequences.

A) Results for 9 iterations of the simulation with different pseudo-random stimulus sequences for an artificial observer with aftereffect length 36 and noise weight 0.33 (as in Figure 1A). Each sequence consists of the same number of stimuli that were presented to single observers in Chopin and Mamassian's study. The short-term aftereffect (blue) occurs consistently (because it is implemented in the artificial observer). Positive correlations occur for variable window lags and durations.

B) Mean correlations results of 200 iterations as in **A**. On average, a simulated observer with only a short-term aftereffect still produces responses that are positively correlated for long window durations and lags.

C) Null distribution of correlation coefficients for some selected window durations and lags. Distributions of *r* values (and Fisher *z'* values, which are normally distributed) are positively biased for large window durations and lags. To statistically assess whether correlations of human observers' responses with long-term stimulus history are higher than expected from just a short-term aftereffect, one should compare empirical correlations at each window position to corresponding null distributions, like those in panel **C**.