Hue shifts produced by temporal asymmetries in chromatic signals depend on the alignment of the first and second harmonics

Andrew Stockman

UCL Institute of Ophthalmology, University College London, London, United Kingdom

G. Bruce Henning

UCL Institute of Ophthalmology, University College London, London, United Kingdom

Peter West

UCL Institute of Ophthalmology, University College London, London, United Kingdom

Andrew T. Rider

UCL Institute of Ophthalmology, University College London, London, United Kingdom

Caterina Ripamonti

UCL Institute of Ophthalmology, University College London, London, United Kingdom

When M- or L-cone-isolating sawtooth waveforms flicker at frequencies between 4 and 13.3 Hz, there is a mean hue shift in the direction of the shallower sawtooth slope. Here, we investigate how this shift depends on the alignment of the first and second harmonics of sawtooth-like waveforms. Below 4 Hz, observers can follow hue variations caused by both harmonics, and reliably match reddish and greenish excursions. At higher frequencies, however, the hue variations appear as chromatic flicker superimposed on a steady light, the mean hue of which varies with second-harmonic alignment. Observers can match this mean hue against a variable-duty-cycle rectangular waveform and, separately, set the alignment at which the mean hue flips between reddish and greenish. The maximum hue shifts were approximately frequency independent and occurred when the peaks or troughs of the first and second harmonics roughly aligned at the visual input—consistent with the hue shift's being caused by an early instantaneous nonlinearity that saturates larger hue excursions. These predictions, however, ignore phase delays introduced within the chromatic pathway between its input and the nonlinearity that produces the hue shifts. If the nonlinearity follows the substantial filtering implied by the chromatic temporal contrast-sensitivity function, phase delays will alter the alignment of the first and second harmonics such that at the nonlinearity, the waveforms that produce the maximum hue shifts might well be those with the largest differences in rising and falling slopes—consistent with the hue shift's being caused by a central nonlinearity that limits the rate of hue change.

Introduction

Mechanisms that signal hue variation are less sensitive to high temporal frequencies than mechanisms that signal luminance or brightness variation (e.g., King-Smith, 1975; King-Smith & Carden, 1976; Noorlander, Heuts, & Koenderink, 1981). This difference is apparent in temporal contrast-sensitivity functions (TCSFs) that show how sensitivity to sinusoidal flicker varies with frequency. Chromatic TCSFs obtained with equiluminant stimuli are generally low-pass in form and fall more steeply with increasing frequency than achromatic or luminance TCSFs, which are band-pass in form and extend to higher frequencies than chromatic TCSFs (e.g., de Lange, 1958; Kelly & van Norren, 1977; King-Smith & Carden, 1976; Metha & Mullen, 1996; Regan & Tyler, 1971; Smith, Bowen, & Pokorny, 1984; Sternheim, Stromeyer, & Khoo, 1979; Stockman, MacLeod, & DePriest, 1991; Tolhurst, 1977). TCSFs are often taken to indicate the attenua-
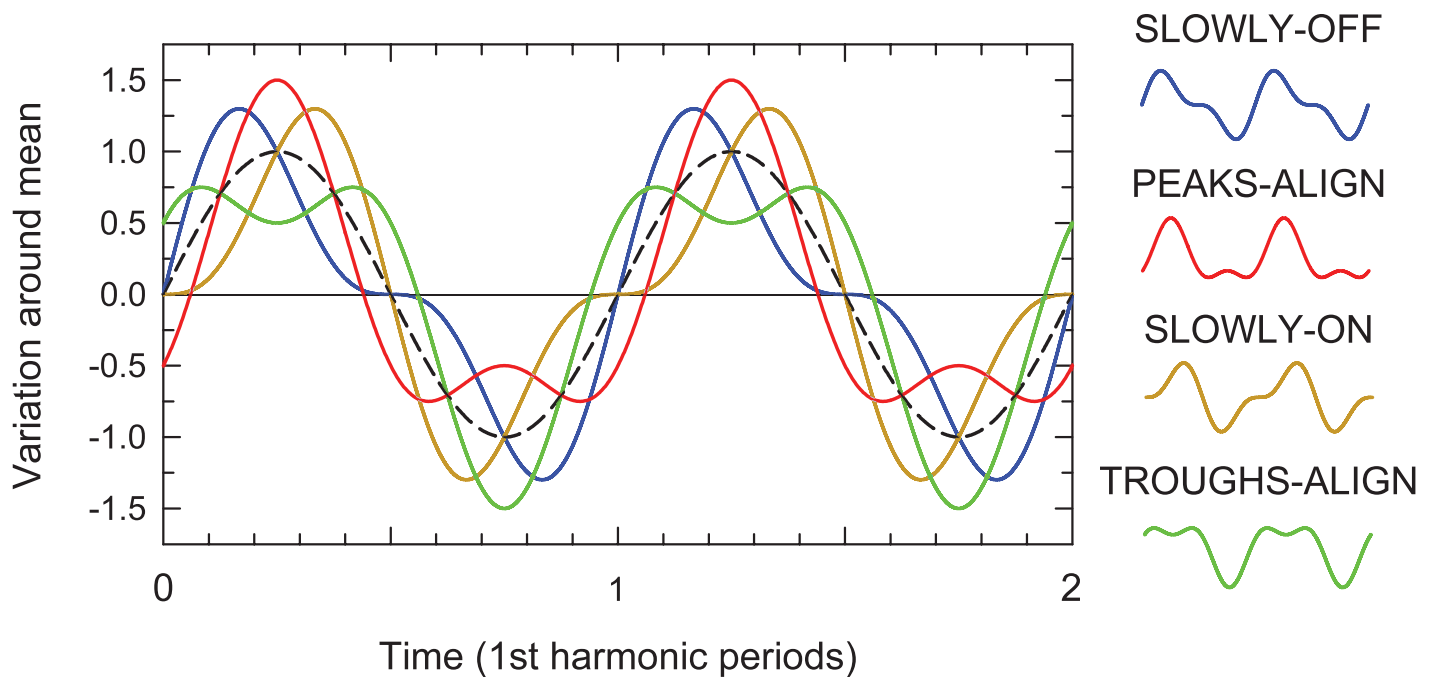
Figure 1. Two cycles of five example waveforms used in the experiments. The dashed black line shows a sinusoid (first harmonic) of period 1 and an amplitude of 1. The other waveforms are made up of the first harmonic (fundamental) with an amplitude of 1 and a second harmonic with an amplitude of 0.5. All the waveforms modulate around a mean of 0. The different lines represent different phases of the second harmonic (in 90° steps from 0° to 270°). We generally refer to the second-harmonic phases as *phase delays* relative to the slowly-off phase of 0°. Thus, the four two-component waveforms, which are shown individually on the right with the names we use to describe them, are slowly-off (0° second-harmonic phase delay, blue line), peaks-align (90° phase delay, red line), slowly-on (180° phase delay, brown line), and troughs-align (270° phase delay, green line).

tion characteristic of temporal filtering in the visual system and are modeled as a cascade of filters. The relative loss of high-frequency chromatic sensitivity is typically attributed to extra filtering stages or to stages with longer time constants, both of which increase the sluggishness of the chromatic system (e.g., Lee, Sun, & Zucchini, 2007; Yeh, Lee, & Kremers, 1995).

Petrova, Henning, and Stockman (2013) and Stockman, Petrova, and Henning (2014) refined the chromatic-filter model by using the distortion in contrast-modulated flicker to measure separately the attenuation characteristics of early and late filters in the chromatic pathway. They found that the early filter acts as a band-pass filter peaking at 10–15 Hz, while the late filter acts as a two-stage low-pass filter that begins to attenuate chromatic flicker above about 3 Hz. Much of the additional loss of high-frequency sensitivity in the chromatic system (when compared to the achromatic or luminance system) appears to be due to the late, low-pass filter. As part of their work, Stockman et al. (2014) identified two nonlinear stages in the chromatic pathway: an early expansive nonlinearity, linked to the half-wave rectification that occurs when signals are segregated into ON and OFF pathways, and a central compressive nonlinearity that follows the late low-pass filter.

In a companion article, we recently reported and characterized a new color phenomenon (Stockman et

al., 2017) in which asymmetries in sawtooth-shaped, flickering L-cone- or M-cone-isolating waveforms produce shifts in the mean hue of the waveforms. Between about 4 and 13 Hz, slowly-on/rapidly-off or slowly-off/rapidly-on sawtooth waveforms that are modulated around a mean yellow-appearing chromaticity exhibit shifts in mean hue in the direction of the slowly changing part of the sawtooth: toward red for slowly-on L-cone and slowly-off M-cone waveforms, and toward green for slowly-off L-cone and slowly-on M-cone waveforms. Measurements of the dependence of the mean hue shifts on modulation depth and on the slope of the sawtooth revealed that the discriminability of the hue changes depended mainly on the second harmonic of the sawtooth. Consequently, in this article only the first and second harmonics of the sawtooth waveforms are used.

The first two harmonics of sawtooth waveforms are given by

$$y(t) = \sin(2\pi f t) + \frac{1}{2}\sin(2\pi 2 f t + \theta), \quad (1)$$

where f is frequency (Hz), t is time (s), $\theta = 0°$ for a slowly-off waveform, and $\theta = 180°$ (π radians) for a slowly-on waveform. The slowly-off and slowly-on waveforms are shown in Figure 1 as the blue and brown lines, respectively. For both waveforms, the second

harmonic has half the amplitude of the first harmonic, so that the amplitude spectra of the two waveforms are identical. The difference between the waveforms lies in the phase of their second harmonics, which differ by 180°. In our experiments, we vary θ.

The main panel of Figure 1 shows two cycles of five waveforms each with a period of 1 second: The dashed black line is the fundamental component alone (the sinusoidal first harmonic), and its phase is fixed at 0°. The other four waveforms (colored lines) are each composed of the first and second harmonics with amplitudes of 1 and 0.5, respectively, but each composite waveform has a different second-harmonic phase. The icons and labels on the right illustrate these composite waveforms and give the names we use to describe them: slowly-off (blue line), peaks-align (red line), slowly-on (brown line), and troughs-align (green line). Each waveform is generated by a different second-harmonic phase.

From Equation 1 it can be seen that positive values of phase (θ) slide the second harmonic to the left (i.e., advance it), while negative values slide it to the right (i.e., delay it). We shall call negative values of phase *phase delays* and will often specify the phase delay of the second harmonic relative to the phase of the first harmonic. (This simply means finding a time when the first harmonic is in sine phase—crossing zero in the positive-going direction as in Figure 1—calling that time 0, and noting the resulting phase θ of the second harmonic, because the waveform can then be described as in Equation 1.) Several features of these waveforms are worth noting: The slowly-off (blue) and slowly-on (brown) waveforms have the greatest asymmetry between their rising and falling sides, while the peaks-align (red) and troughs-align (green) waveforms have the greatest asymmetry between their excursions above and below the mean.

In this article, we made hue matches as a function of the phase of the second harmonic using waveforms with fundamental frequencies ranging from 0.63 to 13.33 Hz (and thus second-harmonic frequencies from 1.34 to 26.67 Hz). Subjectively, and in terms of the matching techniques that could be used, this frequency range split into two distinct regions: the low-frequency region from 0.63 to 2 Hz in which the hue variations produced by the second harmonic can be followed by the observer, and the region from 4 to 13.33 Hz in which observers report the hue variations as chromatic flicker superimposed on a steady light, the hue of which varies with second-harmonic phase.

General methods

This research adhered to the tenets of the Declaration of Helsinki.

Observers

One female observer (KR) and five male observers (AR, AS, MG, RTE, and VL) participated; KR (Caterina "Katia" Ripamonti), AR, and AS are authors. They were all experienced psychophysical observers with normal color vision and normal spatial acuity, except AR, who has mild deuteranomaly and therefore performed only M-cone-isolating experiments.

Testing system

We used a calibrated 21-in. Sony FD Trinitron CRT and a VSG2/5 Visual Stimulus Generator (Cambridge Research Systems Ltd., Rochester, UK) that together provide an intensity resolution of up to 14 bits per gun. A ColorCAL colorimeter (Cambridge Research Systems Ltd.) was used to measure the luminance of each phosphor for linearization using lookup tables and for daily calibration of the experimental conditions. Each of the red, green, and blue guns of the monitor were individually linearized. A Radoma spectroradiometer (Gamma Scientific, La Jolla, CA) was used to measure the spectral power distributions of each of the three CRT primaries.

The refresh rate of the monitor was set to 160 Hz (or to 100 Hz for the 0.63-Hz stimulus frequency) with a spatial resolution of 800 × 600 pixels. A six-key response box was used to collect the observers' responses. The timing of stimulus generation was implemented by the VSG system and was independent of the operating system on the host PC.

Stimuli

Observers sat 1 m from the computer monitor. From that distance, the stimuli presented on the CRT were semicircular half-fields, whose diameter subtended 5.7° at the observer's eye, separated by a vertical gap of 0.6° visual angle. The stimuli were viewed binocularly through natural pupils. Fixation was not constrained, and observers were free to fixate either half of the split field in making their judgments. The mean overall luminance was 42.39 cd/m^2.

The half-fields were temporally modulated (flickered). Simple waveforms, such as sawtooth waveforms or variable-duty-cycle rectangular waves, were calculated in real time; other waveforms, such as combinations of sinusoids, were precalculated, sampled at the appropriate frame rate, and displayed on the CRT.

The frame rate of the CRT of 160 Hz (at stimulus frequencies > 0.63 Hz) places the Nyquist frequency, above which sinusoids will produce aliasing, at 80 Hz.

Thus, in the frequency range of interest for mean hue shifts (up to 13.33 Hz), the second harmonic of the two-component stimuli remained well below the Nyquist frequency. (The square waves used in the hue-matching task were restricted to the first, third, and fifth harmonics so as not to exceed the Nyquist frequency.) The exceptions were the variable-duty-cycle rectangular waveforms for which the sixth harmonic and above exceed the Nyquist frequency. However, over the range of duty cycles that we employed from (about 25% to 75%), those harmonics are small, and control experiments in which the higher harmonics of the waveforms were removed showed that they affected neither the hue appearance nor the judgments.

The frequencies of the stimuli were chosen to have an integer number of frames in a period. Note that any change in the stimulus requires at least one frame to occur; thus, the rapid ramp of sawtooth waveforms at or above 1 Hz was not instantaneous, but took 6.25 ms (10 ms for the 0.63-Hz flicker).

Cone modulations in each half-field were around a yellow point (CIE $x = 0.436$, $y = 0.489$) with a luminance of 42.39 cd/m^2. The pupil size was not controlled, but with a 3-mm diameter pupil—appropriate for our luminance (Watson & Yellott, 2012)—the target would have had a mean retinal illuminance of 2.18 log$_{10}$ photopic trolands. The phosphors at their maximum intensities were red: $x = 0.615$, $y = 0.343$, 28.33 cd/m^2; green: $x = 0.283$, $y = 0.611$, 60.26 cd/m^2; and blue: $x = 0.154$, $y = 0.079$, 10.43 cd/m^2.

Standard methods were used for generating cone-isolating stimuli: The spectral power distributions of the three phosphors were multiplied by the Stockman and Sharpe (2000) cone fundamentals and the phosphor combinations that produced silent substitutions for each cone type calculated (Estévez & Spekreijse, 1982). The main waveforms of interest in these experiments were composed of two sinusoids—the fundamental or first harmonic and the second harmonic at twice the frequency and half the amplitude of the first. The phase of the second harmonic was varied. The M- or L-cone modulations of the first harmonic of the two-component waveforms were fixed at 10% so that the modulation of the second harmonic was always 5%.

In the first experiment, observers matched the time-varying (AC) hue of the two-component flickering waveforms by varying the modulation of a square-wave flicker of the same fundamental frequency.

Slew-rate-limited and saturation model simulations

Because the slew-rate-limited and saturation models are inherently nonlinear, we cannot employ standard linear-systems techniques alone to understand and predict their behavior. Therefore, we have used simulations implemented in MATLAB and Simulink (MathWorks, Natick, MA). These simulations are done by generating the waveforms using standard MATLAB functions (e.g., square, sin, sawtooth) and then passing the waveforms through the Simulink Saturation or Rate Limiter function blocks. The output can then be viewed using the Scope function block or analyzed using either Simulink function blocks or MATLAB functions. The input and output waveforms shown in Figures 12 and 13 are taken directly from MATLAB.

Curve fitting

Curve fitting was used to obtain sinusoidal fits to some of the data. The fits used the standard nonlinear Marquardt–Levenberg curve-fitting algorithm (Levenberg, 1944; Marquardt, 1963) implemented in Sigmaplot (Systat, San Jose, CA) to minimize the sum of the squared differences between the data and model predictions. Note that for nonlinear regression, the R^2 goodness of fit is an inappropriate metric (e.g., Kvalseth, 1985; Spiess & Neumeyer, 2010), so instead we give the standard error of the regression, which can be used to assess the precision of the predictions.

Results

Experiment I: Red–green hue matching at low temporal frequencies

Introduction

In Experiment I, we combined first and second harmonics to produce a compound waveform in which the second harmonic had half the modulation of the first (i.e., in the same ratio as for sawtooth waveforms). We then investigated how the hue appearance of the compound waveform depended on the phase of the second harmonic. Up to about 4 Hz, we found that the hue appearance produced by L- or M-cone flicker is seen mainly as a hue alternation between reddish and greenish, but with some intermediate yellow hues appearing on the shallower slopes of some waveforms. The reddish and greenish hues were perceptually separated in the sense that observers were able to attend to either the redder or the greener hue, but not both at the same time. This separation lent itself to a relatively simple matching procedure in which, in one set of runs, observers used square-wave flicker of variable modulation to match the most redward hue variations produced by a compound waveform of the
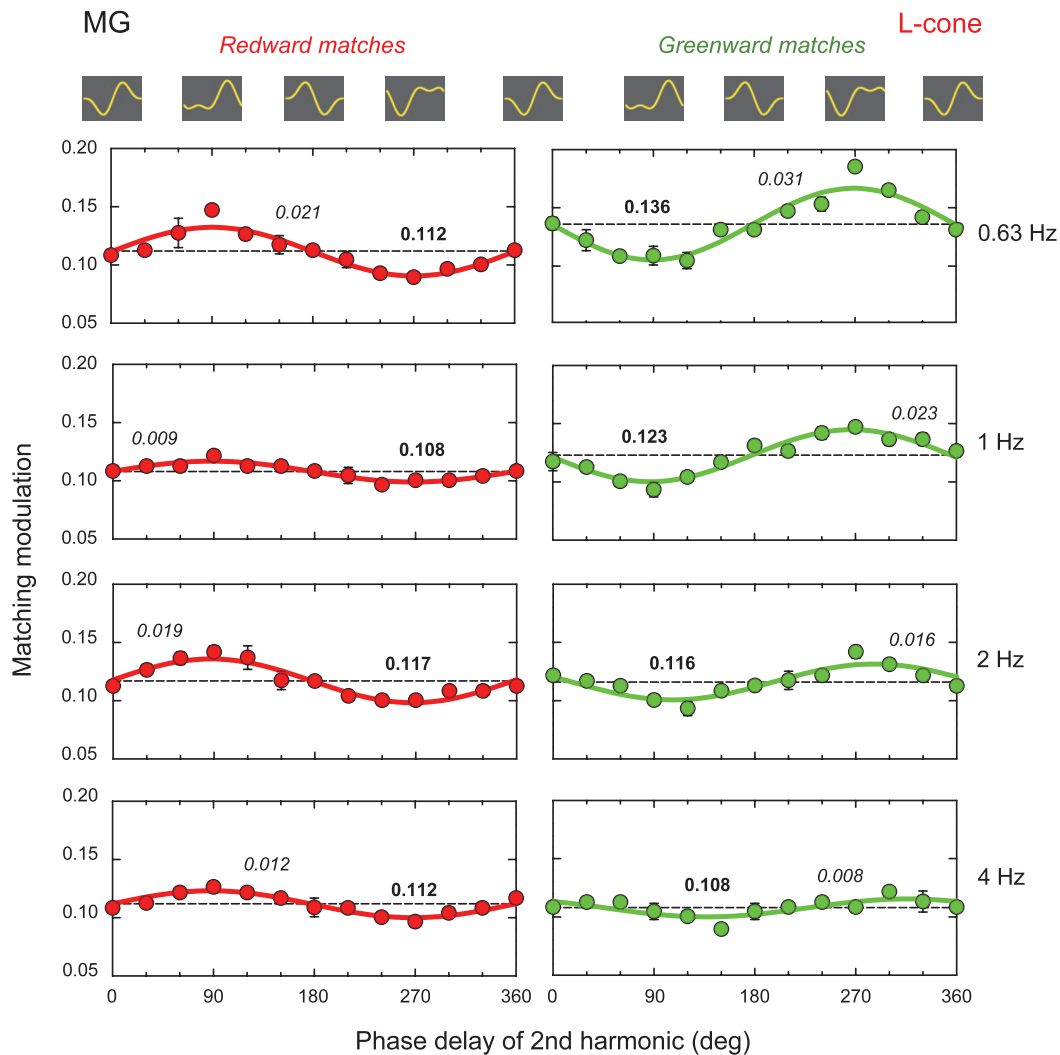
Figure 2. L-cone square-wave modulations for MG that matched the redward excursions (left-hand panels, red circles) and greenward excursions (right-hand panels, green circles) of the two-component L-cone waveforms plotted as a function of the phase delay of the second harmonic (degrees). 0° indicates that the second harmonic, like the first harmonic or fundamental, is in sine phase. Error bars are ±1 standard error of the mean. The red and green lines are single periods of a sinusoid that best fits (least-squares) the redward and greenward matches, respectively. The best-fitting sinusoidal mean is plotted as the dashed horizontal line and noted in bold numerals in each panel. The best-fitting sinusoidal amplitude is noted in italic numerals. The panels show fundamental frequencies ranging from 0.63 Hz at the top to 4 Hz at the bottom. The icons above the panels illustrate the rapid-on, peaks-align, rapid-off, and troughs-align waveforms at phase delays of 0°, 90°, 180°, and 270°, respectively. (The waveforms at 0° and 360° are, of course, identical.)

same fundamental frequency but fixed modulation; and, in a different set of runs, they matched the most greenward hue variations.

Methods

The stimuli were presented in two semicircular half-fields. One half-field contained a two-component waveform with the first harmonic having 10% cone modulation and the second having 5% cone modulation. The other half-field contained square-wave flicker of the same fundamental frequency as the two-component waveform. The square wave was composed of only its first, third, and fifth harmonics, with modulations in the proportions of 1.00:0.33:0.20; and the matching square-wave modulations given in the following are the modulation of the first harmonic of the square wave. The observers varied the square-wave modulation in logarithmic steps of 0.05 in order to match either the maximum redward or the maximum greenward excursions of the two-component stimulus. (The cone modulation of the matching square wave, although variable, could not exceed 25% because of the limitations of the system.) The left and right positions
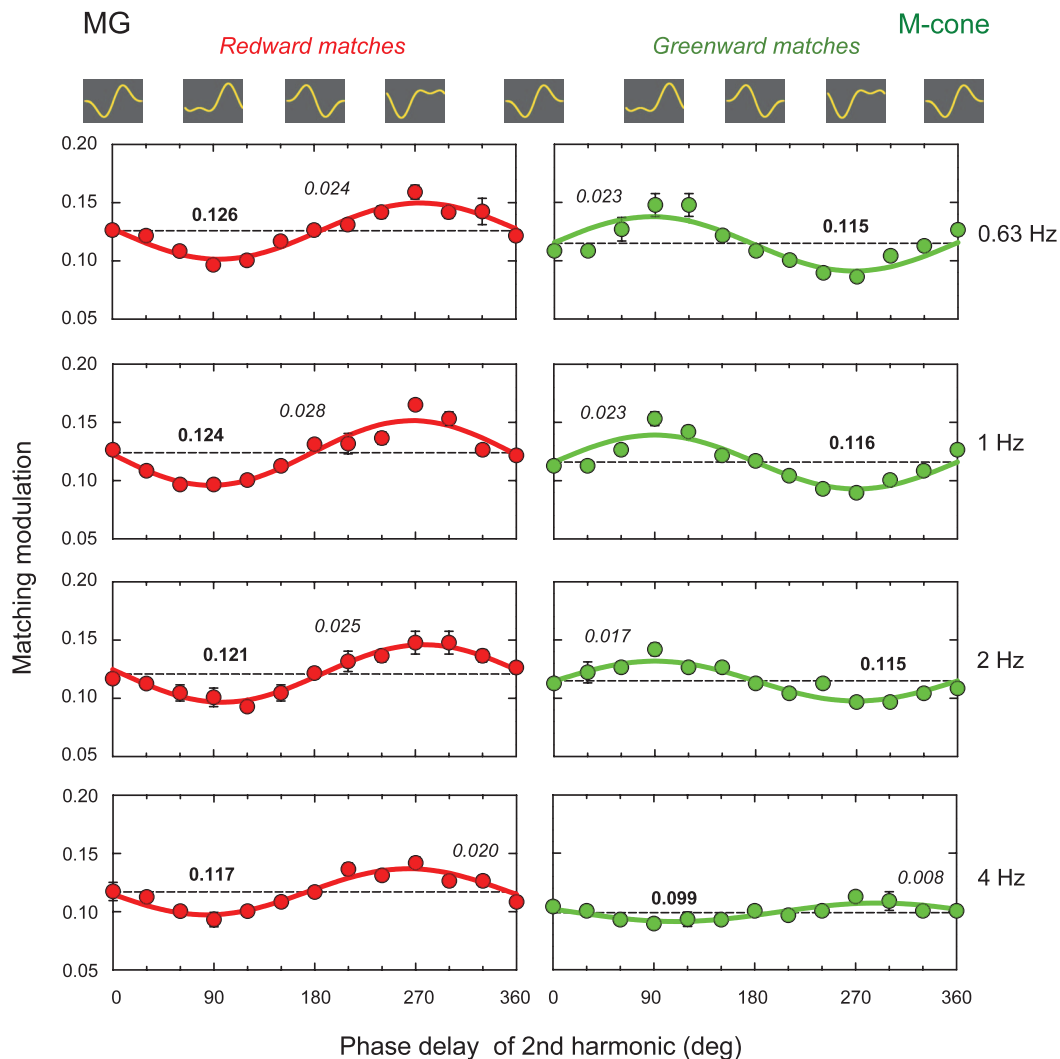
Figure 3. Same as Figure 2 but for M-cone-isolating stimuli.

of the two waveforms was randomized. Observers could also delay or advance the square wave in time in order to align its perceptual extrema with those of the two-component waveform; the observers found the matching task easier when the redward or greenward excursions of both stimuli appeared roughly synchronous. The temporal frequencies of the two-component and square-wave flicker were varied together.

Redward and greenward matches were done separately, and the method of adjustment was used. The matches were repeated three times and then averaged. The phase of the second harmonic of the two-component waveform was varied in sequence within a run from 0° to 360° in 12 steps of 30°. Since the first and last phases produce the same waveform, the repeat measurement at the beginning and end of a run allowed us to check for any long-term changes in the match—we found no evidence for such changes.

Results

Figure 2 shows the matches made by MG for two-component waveforms that drive the L-cones. The left-hand column of panels shows the redward matches, and the right-hand column, the greenward matches. In both columns, the fundamental frequency increases from 0.63 Hz in the top panel to 4 Hz in the bottom panel. The error bars are ±1 standard error of the mean taken over the three runs. Each panel shows the square-wave modulation (plotted as the modulation of the fundamental component of the matching square wave) that produced the match as a function of the phase delay of the second harmonic; in terms of Equation 1, the phase delay in degrees is $-\theta$. The icons at the top illustrate, and are aligned with, the waveform shapes at second-harmonic phase delays of 0°, 90°, 180°, 270°, and 360°.

Both redward and greenward matches could be achieved by adjusting the modulation of the square wave. When matched, the fundamental component of
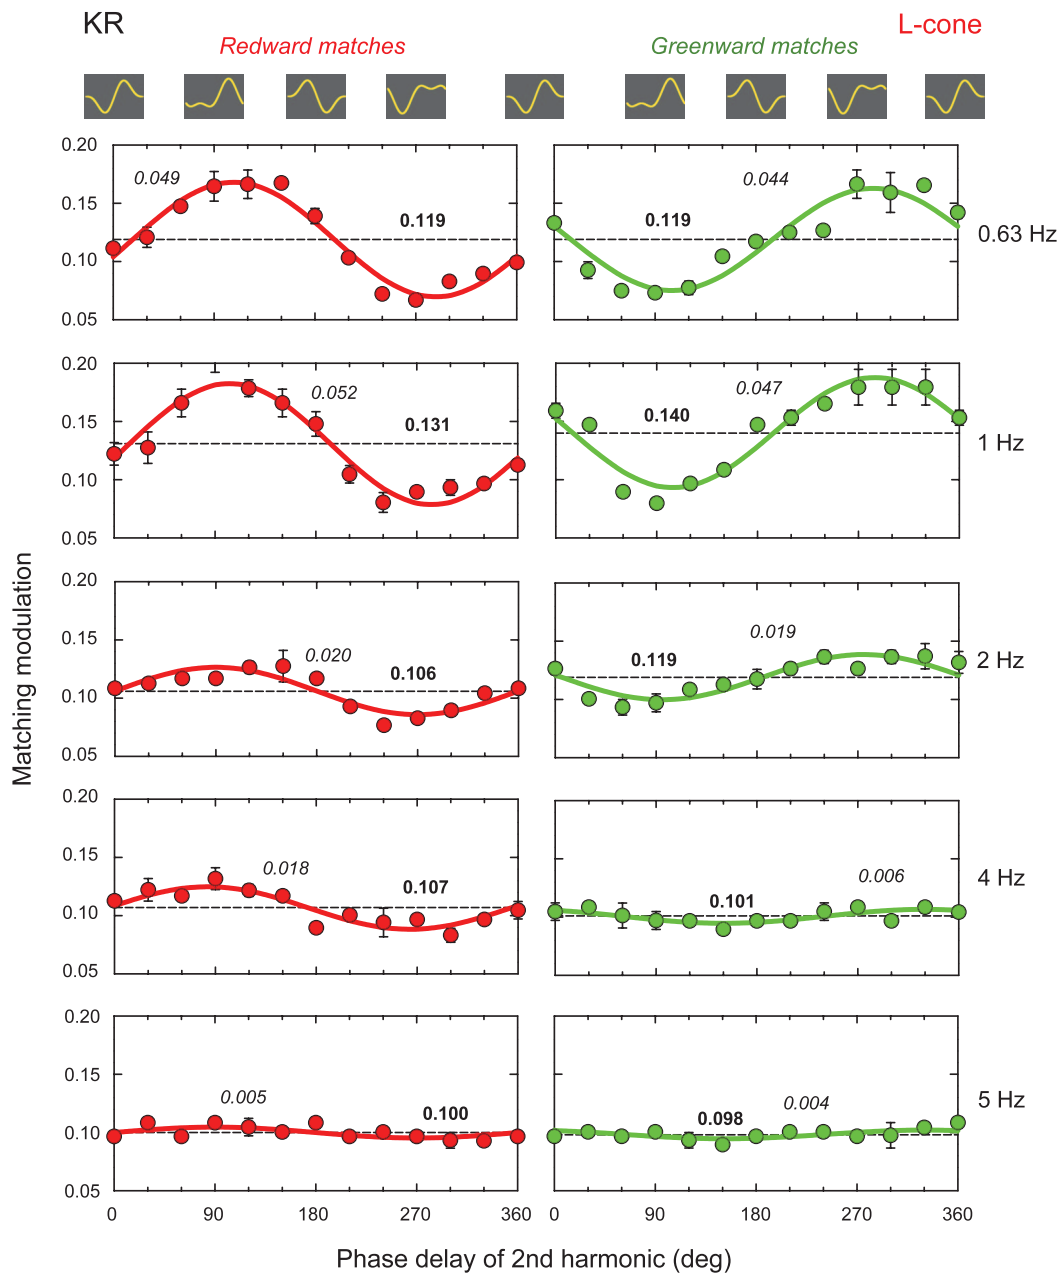
Figure 4. Same as Figure 2 but for observer KR, who also made matches at 5 Hz (bottom panels).

the square (plotted along the ordinate) varied in a way that appears to depend sinusoidally on the phase of the second harmonic. The solid red and green lines show the sinusoid that best fits (least-squares) the data. The best fits to the redward and greenward matches are in opposite phase. Thus, at the second-harmonic phase delay where the redward match reaches a maximum, the greenward match reaches a minimum, and vice versa. By and large, the variation of the matching modulation with second-harmonic phase declines with increasing frequency, dropping to about 0.01 at 4 Hz. Between 0.63 and 4 Hz, the second-harmonic phase delay for L-cone modulation that yielded maximal perceived redness (and minimal perceived greenness)

was near 90° (corresponding to the peaks-align waveform), while a phase delay near 270° (the troughs-align waveform) minimized redward and maximized greenward excursions. By contrast, the second-harmonic phase delay for M-cone modulation that yielded maximal perceived greenness (and minimal perceived redness) was near 90° (corresponding to the peaks-align waveform), while a phase delay near 270° (the troughs-align waveform) minimized greenward and maximized redward excursions. This is much as expected if the perceptually alternating colors are mainly determined by the extremes of the stimulus waveform.

The solid colored lines through the data in each panel are offset sinusoids that best fit each set of data.
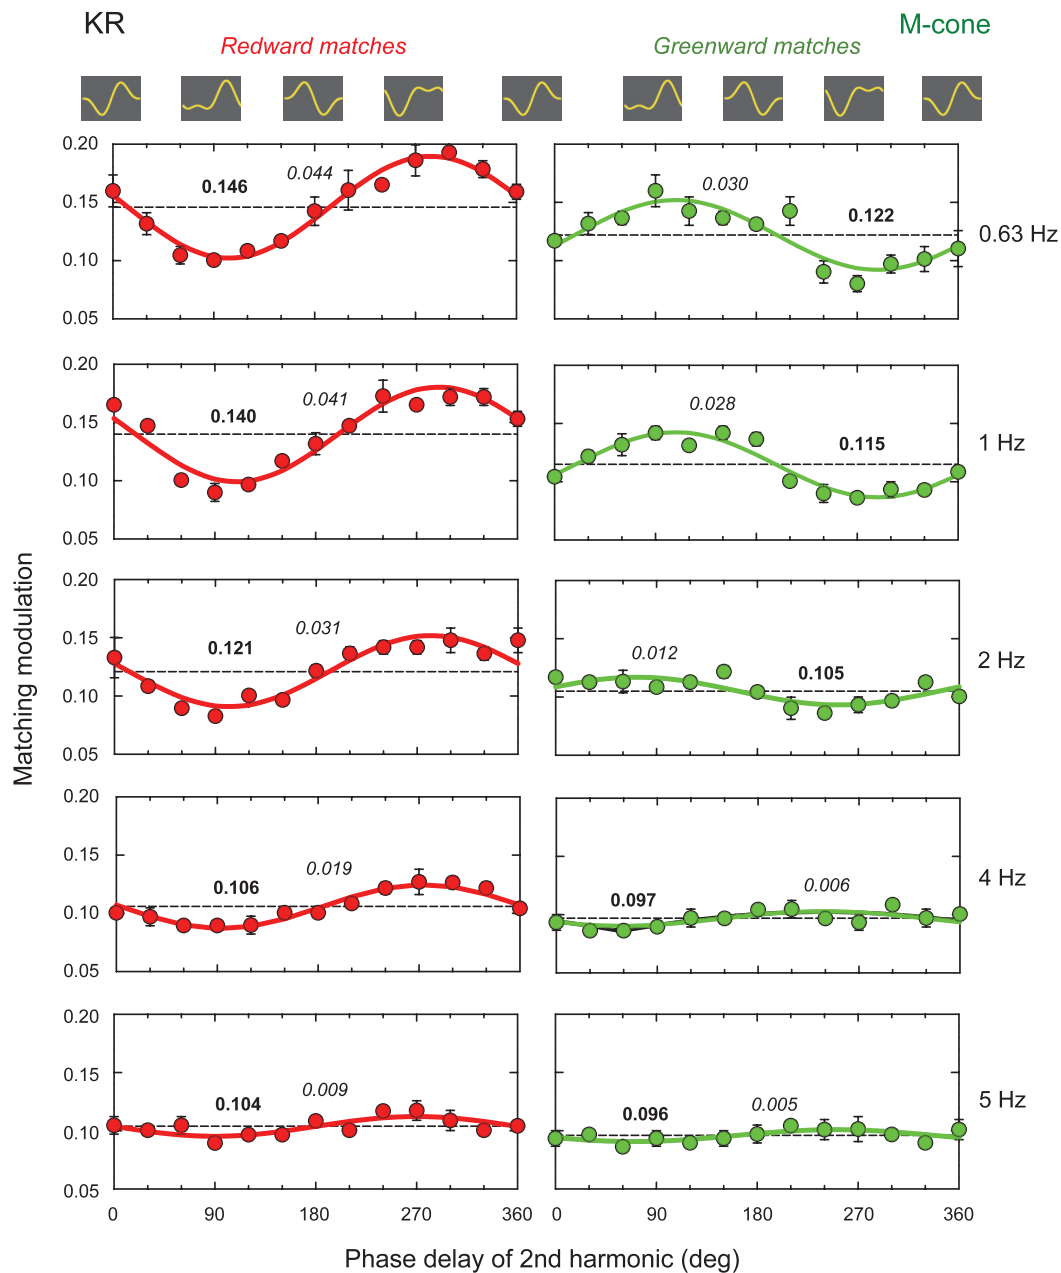
Figure 5. Same as Figure 3 but for observer KR, who also made matches at 5 Hz (bottom panels).

The offset is the constant or mean value around which the best-fitting sinusoid varied (indicated by the dashed horizontal line). This corresponds to the mean matching modulation of the fundamental of the square wave. The numbers in each panel give the mean (in italic) and amplitude (bold) of the best-fitting (least-squares) sinusoid. (The means, amplitudes, and phases of the offset sinusoids in Figures 2 through 5 are plotted in Figure 6.)

Figure 3 is in the same form as Figure 2 and shows MG's square-wave matches for the two-component waveforms that drive the M-cones. The M-cone matches are generally similar to the L-cone matches

with one major difference: The sinusoidal fits to the data are shifted from the fits to the L-cone data by approximately 180°. This 180° shift reflects the opponency of the L- and M-cone signals in producing hue. Thus, between 0.63 and 4 Hz for MG's M-cone redward matches, and between 0.63 and 2 Hz for the greenward matches, the maximum redward and minimum greenward excursions occur near the 270° phase delay that corresponds to the troughs-align waveform, while the minimum redward excursions and maximum greenward excursions occur at second-harmonic phase delays near the 90° phase delay that corresponds to the peaks-align waveform. (The sole exception to this

	L-cone		M-cone	
	Redward	Greenward	Redward	Greenward
MG				
0.63 Hz	0.0061	0.0086	0.0056	0.0106
1 Hz	0.0024	0.0054	0.0076	0.0087
2 Hz	0.0050	0.0063	0.0042	0.0056
4 Hz	0.0033	0.0056	0.0052	0.0036
KR				
0.63 Hz	0.0088	0.0133	0.0063	0.0116
1 Hz	0.0114	0.0122	0.0102	0.0070
2 Hz	0.0079	0.0080	0.0095	0.0078
4 Hz	0.0076	0.0043	0.0043	0.0057
5 Hz	0.0045	0.0042	0.0056	0.0043

Table 1. Standard errors of sinusoidal fits to the square-wave red–green hue matches for L-cone and M-cone modulations in the redward and greenward hue directions for observers MG and KR at fundamental frequencies from 0.63 to 4 Hz.

pattern is that, for MG, the very-low-amplitude greenward matches at 4 Hz appear to be of the same phase as the redward ones.)

The square-wave matches for the L-cone and M-cone two-component waveforms for KR are shown in Figures 4 and 5, respectively. KR also made matches at 5 Hz. In general, the data for KR are similar to those for MG. Again, the phase-dependent effect of the second harmonic falls rapidly with increasing frequency.

Both observers tried to make matches at 6.67 and 8 Hz, but the matches were unsatisfactory. The matches break down above 4 Hz partly because, at some phases, the two-component waveforms no longer vary in hue between red and green, but are sometimes limited to intermediate hues. In addition, the variations are seen as chromatic flicker superimposed on a steady light, the hue of which varies with second-harmonic phase. At most phases, these complex hue variations cannot be matched against a simple square wave, but—as we show in Experiment III—they can be matched against a rectangular waveform with a variable duty cycle.

Discussion

The relation between the observers' redward and greenward square-wave hue matches and the second-harmonic phase delay is well described by an offset sinusoid. Table 1 gives the standard errors for these sinusoidal fits. The standard errors are much smaller than the variation with second-harmonic phase at lower frequencies, but the fits are less secure at 4 and 5 Hz. The red and green solid lines plotted through the data in each panel of Figures 2 through 5 show the best-fitting offset sinusoid. In each case, we used a

least-squares criterion to find the best-fitting offset (mean), amplitude, and phase. Overall, the three parameter fits are remarkably good and characterize the matching data well. The mean of the fit provides an estimate of the size of the hue variation averaged over all phases of the second harmonic (and corresponds to the mean size of the matching first harmonic), the amplitude of the fit provides an estimate of the size of the second-harmonic-dependent hue variation, and the phase of the best fit allows us to locate the phase delays of the second harmonic that produce the maximum and minimum redward or greenward hue excursions.

Figure 6 plots each of the three parameters of the best-fitting sinusoids as a function of fundamental frequency for the L-cone matches (left-hand panels) and for the M-cone matches (right-hand panels). Data for KR are shown as circles, those for MG as triangles; the redward matches are plotted as red symbols and the greenward matches as green symbols. The top panels show the best-fitting means (offsets), the middle panels show the best-fitting amplitudes (reflecting the influence of second-harmonic phase), and the bottom panels shows the best-fitting phases. The error bars in Figure 6, if larger than the symbols, show ±1 standard error for the fit of each parameter.

The best-fitting means (top panels) decrease with frequency and approach the 10% modulation of the fundamental component of the two-component wave-form at the higher fundamental frequencies. If the hue variation at the fundamental frequency depended on just the first harmonic, the mean hue shift should align with the dashed horizontal lines in the upper panels, which correspond to the 10% modulation of the first harmonic. Above about 4 Hz, the second harmonic has little effect on the matches, and observers simply match the modulation of the fundamental of the square wave to that of the fundamental of the two-component waveform. At lower frequencies, the second harmonic adds to the hue variation produced by the first harmonic. Below 4 Hz the hue variation exceeds 10%, indicating that the second harmonic increases the hue variation produced at the fundamental frequency. The loss of influence of the second harmonic above a fundamental frequency of 4 Hz in these matches suggests that the variation in the waveform produced by the second harmonic is attenuated by a filter in the chromatic pathway with a cut-off as low as 3 or 4 Hz. Such a filter (with a cutoff frequency of 3.15 Hz) was identified as the late filter in the chromatic pathway by Petrova et al. (2013) in their use of flicker distortion to dissect the chromatic pathway.

The amplitude of the fitted sinusoid (middle panels) quantifies the size of the contribution of the second harmonic to the hue shifts as a function of the fundamental frequency. The additional phase-dependent hue variation caused by the second harmonic is
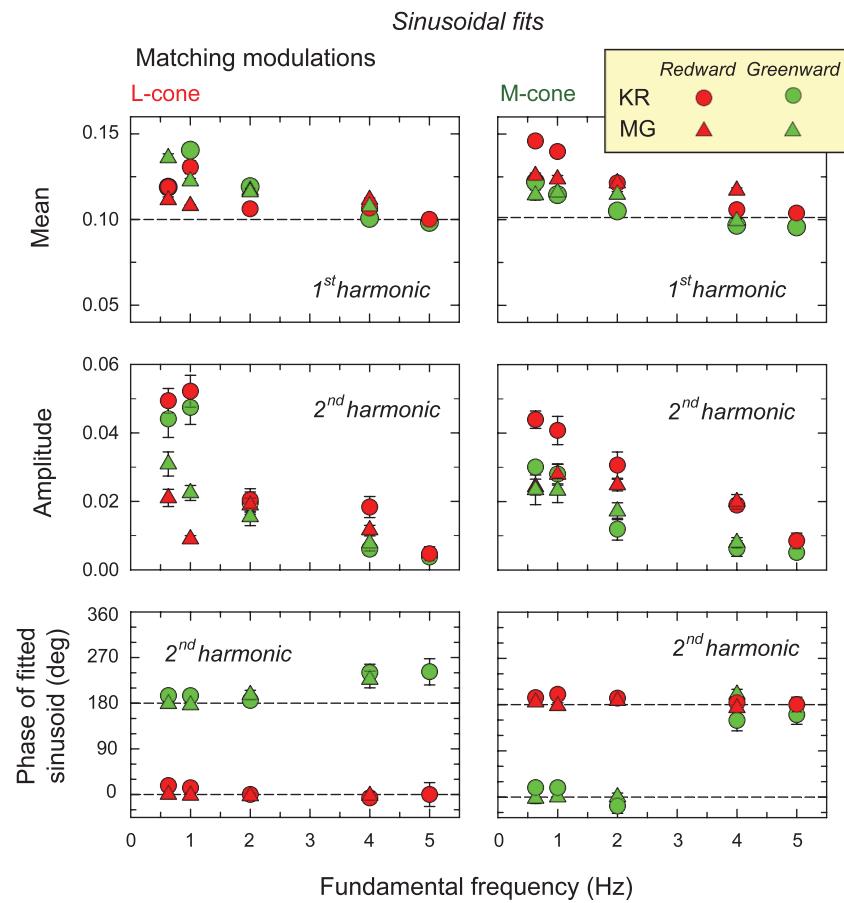
Figure 6. Best-fitting values of the three parameters of the sinusoids fitted to the data in each panel of Figures 2 through 5: means (upper panel), amplitudes (middle panel), and phases (lower panel). The sinusoids were fitted to the redward (red symbols) and greenward (green symbols) square-wave matches for KR (circles) and for MG (triangles). Fits for L-cone data are shown in the left-hand column of panels, and for M-cone data in the right-hand column. The horizontal dashed lines in the upper panel mark the 10% modulation of the first harmonic, and the horizontal lines in the bottom panel emphasize fitted sinusoidal phases of 180° and 0°. The error bars, which are only visible when they are larger than the symbols, show the standard error of the fitted parameter.

largest at 0.63 and 1 Hz but then declines rapidly with increasing frequency and is close to zero at 5 Hz.

The bottom panels plot the phase of the sinusoid that best fits the matching data. The dashed horizontal lines mark the best-fitting phases of 0° and 180°. A phase of 0° in the best-fitting sinusoid means that the maximum and minimum excursions occur at second-harmonic phase delays of 90° (peaks-align) and 270° (troughs-align), respectively, whereas a phase of 180° in the best-fitting sinusoid means that the maximum and minimum excursions occur at second-harmonic phase delays of 270° (troughs-align) and 90° (peaks-align). The phase of the fits, which are roughly independent of frequency, shows that the maximum redness and minimum greenness is perceived when the peaks of the first and second harmonics of the L-cone waveforms align at 90° or the troughs of the first and second harmonics of the M-cone waveforms align at 270°; and the maximum greenness and minimum redness is perceived when the peaks of the first and second

harmonics of the M-cone waveforms align at 90° or the troughs of the first and second harmonics of the L-cone waveforms align at 270°—all of which emphasizes the opponency of red and green hues.

Up to about 4 Hz, the results appear inconsistent with the slew-rate model, since in that model the maximum redward and greenward hue excursions would be expected when the slope of the asymmetric waveforms are shallowest—at second-harmonic phase delays of at 0° and 180° when the composite waveforms are slowly-off or slowly-on, respectively. However, the failure to find maxima at these phase delays may simply mean that at low temporal frequencies, the slew-rate limit does not significantly constrain either the fast or slowly changing regions of the composite waveforms.

The results suggest that observers are able to follow the temporal variations in cone excitation produced by our low-frequency stimuli and so match the maximum redward and greenward hue excursions reasonably well.
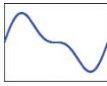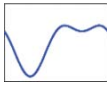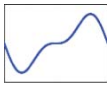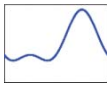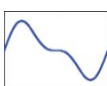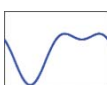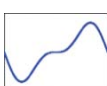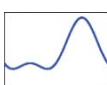
	Waveform	Variable appearance	Mean appearance
L-cone			
Slowly-off		Red and green. Green on longer.	Greenish
Troughs-align		Red/pale orange, achromatic.	Orange
Slowly-on		Red/yellow, no green.	Orange
Peaks-align		Brief red flashes on lime green.	Lime green
M-cone			
Slowly-off		Patches of red on orange background.	Orange
Troughs-align		Brief red flashes on lime green.	Lime green
Slowly-on		Green/red. Symmetric.	Slightly green
Peaks-align		Orange/achromatic.	Orange

Table 2. Phenomenological observations.

Experiment II: Phenomenological hue appearance above 4 Hz

At frequencies above 4 Hz, the hue variations appear as chromatic flicker superimposed on a steady light, the mean hue of which varies with second-harmonic phase. In this brief informal experiment, we describe the subjective appearance of the stimuli at 6.67 and 8 Hz. Observers viewed eight types of M-cone- or L-cone-isolating, two-component, flickering waveforms of the same time-averaged chromaticity and luminance. The modulation of the first harmonic was 10%, and that of the second, 5%. The observers were asked to describe and contrast the appearance of various waveforms, and the descriptions of several normal observers are summarized in Table 2. The column labeled "Waveform" illustrates one cycle of each of the waveforms we used, and the preceding column gives the names we used for the waveforms.

In general, the mean hue shifts (or *DC shifts*, as we sometimes call them) of the peaks-align and slowly-off waveforms were in the same direction, as were the troughs-align and slowly-on waveforms. The chromatic flicker seemed to depend on the peak-to-peak excursion of the waveforms. Thus, slowly-on and slowly-off waveforms appeared to flicker with greater contrast than peaks-align and troughs-align waveforms and produced a symmetrical hue variation around the mean hue. In contrast, the peaks-align and troughs-align waveforms produced more asymmetrical chromatic flicker, with the hue variation in the direction of the greater excursion being larger than that in the opposite direction.

Adding the third harmonic at the amplitude and phases appropriate for sawtooth waveforms to the slowly-on and slowly-off two-component waveforms had no noticeable effect on their appearance.

Experiment III: Red–green hue matching against variable-duty-cycle rectangular waves

Introduction

In Experiment III, we again combined a sinusoidal waveform (first harmonic) with its second harmonic at

half the modulation of the first and varied the phase of the second harmonic. Above 4 Hz, however, we had to adopt a different matching technique because the hue appearance of the two-component waveforms is much more complex than at the lower frequencies used in Experiment I; in addition to seeing chromatic flicker (which we sometimes call the *AC component*) at the fundamental frequency, observers reported a shift in the mean (DC) hue (see Table 2). We experimented with several methods for matching these hues, and eventually discovered that a satisfactory match to both the time-varying and mean hue shifts could be achieved by matching the two-component waveforms against a rectangular-wave stimulus with variable duty cycle and modulation. The duty cycle of a rectangular waveform simply refers to the percentage of a single cycle of the waveform during which the waveform is high rather than low. Thus, a 50% duty-cycle rectangular waveform is a square wave. Initially, we were surprised that this method was so successful, but it later became clear that the matches depended on the similarities between the first- and second-harmonic components of the test and matching stimuli (see Figure 11).

Methods

The stimuli were again presented in 5.7° semicircular half-fields. As before, one stimulus had a first harmonic of 10% cone modulation combined with a second harmonic in some fixed phase and fixed 5% cone modulation. In different conditions, the phase of the second harmonic was chosen to produce waveforms with second-harmonic phase delays ranging from 0° to 360° in 30° steps. The other stimulus was rectangular-wave flicker, in which both the duty cycle and the modulation could be varied by the observer to match the appearance of the two-component stimuli. The left and right positions of the two-component and matching waveforms were randomized. In the experiment, the duty cycle could be varied in 5% steps, and the modulation in logarithmic steps of 0.05. The temporal frequencies of the two-component flicker and the rectangular-wave flicker were varied together.

Changing the duty cycle has a number of effects: The contrast of the fundamental grows as the duty cycle decreases, but the mean cone excitation decreases. We offset the decrease in mean excitation by adding a steady light to keep the mean constant. This manipulation also had the effect of keeping the modulation of the fundamental component roughly constant while allowing the ratios of the amplitudes of the harmonics of the rectangular wave to follow their normal dependence on duty cycle. Adjusting the rectangular stimuli to keep the time-averaged mean amplitude of the waveform constant restricted the range of duty cycles that we could produce, but that did not affect the

matches, which were always within the range of duty cycles available. Although observers were able to adjust the modulation of the rectangular wave, their modulation settings varied relatively little with the phase of the second harmonic and are therefore not shown. Observers could again alter the synchrony of the two waveforms by delaying or advancing the rectangular wave, because they found the task easier when the chromatic flicker excursions appeared synchronous.

Results

After some practice, both observers could satisfactorily match the two-component flicker waveforms by adjusting the modulation and the duty cycle of rectangular-wave flicker. Figures 7 and 8 show the results for two observers separately; the format is comparable to Figures 2 through 5. In Figures 7 and 8, the duty cycle of the rectangular waveform (%) that matches the two-component waveform is shown as a function of the phase delay of the second harmonic.

Figure 7 shows the data for AS, with L-cone matches in the left-hand column and M-cone matches in the right-hand column. The matching duty cycle is shown as a function of the phase delay of the second harmonic (degrees) and both the mean hue shift (DC) and the hue excursions (AC) matched at the given duty cycle. Frequency increases from 4 Hz in the top panels to 13.3 Hz in the bottom panels. The icons at the top of each column illustrate the waveforms at the corresponding second-harmonic phase delays. The icons on the right-hand side show examples of rectangular waveforms with 70% and 30% duty cycles. The duty cycles of the matching waveform reach a maximum (longest positive half-cycle, shortest negative half-cycle) near the 270° phase that corresponds to the troughs-align waveform, whereas they reach a minimum (shortest positive half-cycle, longest negative half-cycle) near the 90° phase that corresponds to the peaks-align waveform, independent of whether L- or M-cone-isolating stimuli are used.

As in Figures 2 through 5, sinusoids that best fit the data are shown as colored lines in each panel. The amplitude and means of the sinusoids that best fit the relation between duty cycle and second-harmonic phase delay in each panel are again indicated in each panel by the italic and bold numerals, respectively. (The phases of the best-fitting sinusoids are plotted as a function of frequency in Figure 9.)

Figure 8 shows the rectangular-wave matches for the two-component L- and M-cone waveforms for KR. The L-cone duty-cycle matches between 4 and 10 Hz are similar to those for AS. The sinusoidal best fits to M-cone duty-cycle matches at 4 and 6.67 Hz are similar to those for AS, but the variations in the adjustments of the duty cycle are much smaller. The L-cone data for
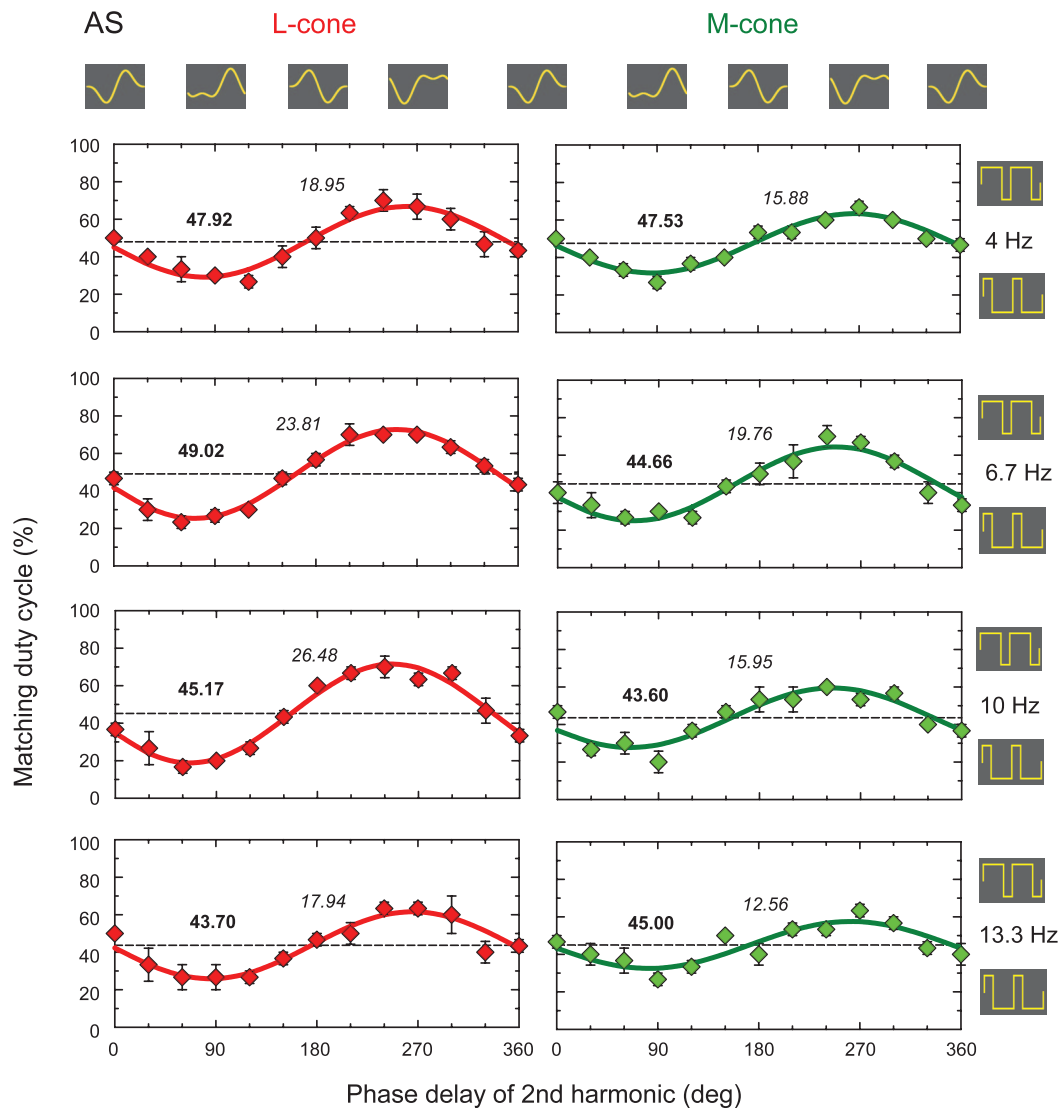
Figure 7. Rectangular-wave duty cycles (%) for AS that matched the appearance of the first- and second-harmonic two-component waveforms as a function of the phase of the second harmonic (degrees). The left-hand column shows matches for L-cone-modulated waveforms (red diamonds) and the right-hand columns matches for M-cone-modulated waveforms (green diamonds) for fundamental frequencies from 4 Hz (top panels) to 13.3 Hz (bottom panels). Error bars are ±1 standard error of the mean. The solid lines are single periods of a sinusoid that best fits the L-cone (red lines) or M-cone (green line) matches. The best-fitting sinusoidal mean is plotted in each panel as the dashed horizontal line and noted in bold numerals. The best-fitting sinusoidal amplitude is noted in italic numerals. The icons above the panels highlight the rapid-on, peaks-align, rapid-off, and troughs-align waveforms at 0°, 90°, 180°, and 270°, respectively.

KR at 13.3 Hz and the M-cone data at 10 and 13.3 Hz show little dependence on second-harmonic phase. (This greater loss with frequency is consistent with work showing that the slew-rate limit for KR is lower than for other observers; see Figure 6 of Stockman et al., 2017).

Phenomenologically, despite maintaining the same physical mean, increasing the duty cycle of the L-cone-modulated rectangular wave above 50% moves the mean hue appearance toward red, while decreasing the duty cycle below 50% moves it toward green. By

contrast, increasing the duty cycle of the M-cone-modulated rectangular wave above 50% moves the mean hue appearance toward green, while decreasing the duty cycle below 50% moves it toward red. Between 4 and 13.3 Hz the largest shifts in mean hue were found for phase delays of about 15° to 30°—close to, but smaller than, the phase delays that correspond to peaks-align and troughs-align waveforms. The AC variations were also similar in appearance when the DC hue shifts matched.
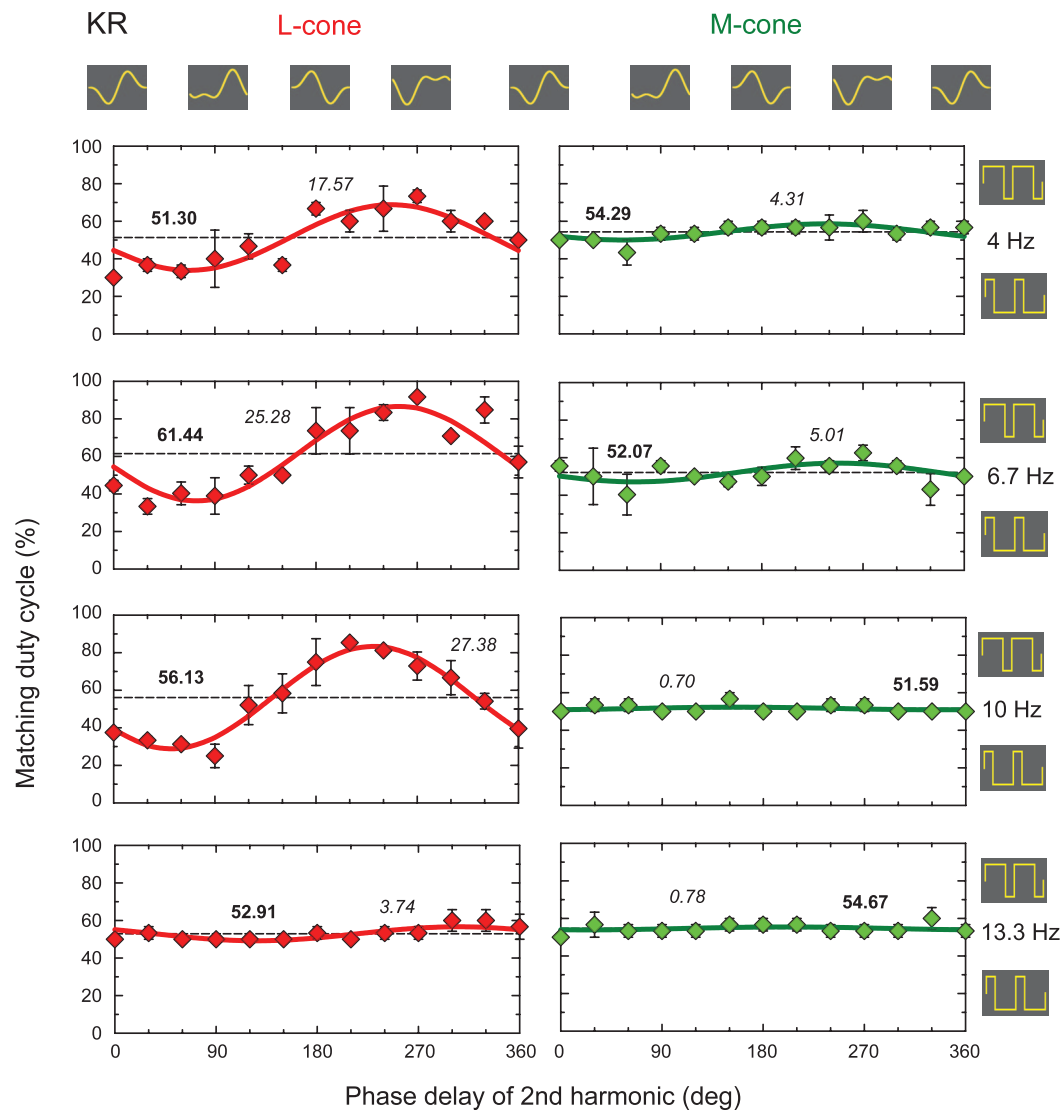
Figure 8. Same as Figure 7 but for observer KR.

Discussion

Like the redward and greenward square-wave matches, the duty-cycle matches for AS and KR when plotted against phase delay are well fitted by single periods of a sinusoid. Table 3 gives the standard errors for these sinusoidal fits. The standard errors are much smaller than the variation with second-harmonic phase for AS for L- and M-cone modulations at all frequencies, and for KR for L-cone modulations from 4 to 10 Hz. The M-cone fits for KR are less secure because the variations in the duty-cycle match with second-harmonic phase are small. The red solid lines aligned with the L-cone matches and the green solid lines aligned with the M-cone matches in Figures 7 and 8 are single periods of sinusoids that were best fits (least-squares) to each set of matching data.

The fitted sinusoids usefully characterize the matching data. The mean of each fitted line provides an

estimate of the mean duty cycle around which the phase-dependent changes vary. The amplitude of the duty-cycle variations gives an indication of the size of the changes that depend on the phase of the second harmonic, and the best-fitting phase allows us to locate the phase delays of the second harmonic that produce the largest DC hue shifts. The best-fitting mean duty cycles (in percentages) and amplitudes are noted in each panel of Figures 7 and 8 as bold and italic numerals, respectively. The best-fitting means are also shown by the dashed horizontal lines in each panel.

The phase of the sinusoids that best fit the data are shown in Figure 9 as a function of frequency: L-cone matches (upper panel) and M-cone matches (lower panels) for AS (orange squares) and KR (blue circles). (The excursions in the duty cycle were too small above 6.7 Hz for KR's M-cone waveforms to usefully estimate the phase of the best-fitting sinusoid.) The error bars in
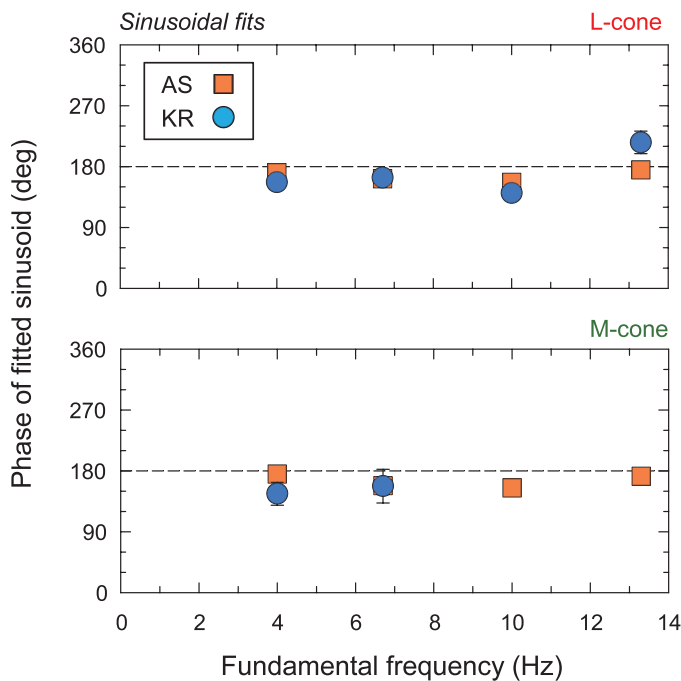
Figure 9. Phases of the sinusoids best fitting the duty-cycle matches in Figure 7 for AS (orange squares) and in Figure 8 for KR (blue circles) for L-cone modulations (upper panel) and M-cone modulations (lower panel) as a function of fundamental frequency. The horizontal dashed lines indicate fitted sinusoidal phases of 180°. The error bars, which are only visible when they are larger than the symbols, show the standard error of the fitted parameter.

Figure 9 show ±1 standard error for the best-fitting phase.

In Figures 7 and 8, the best-fitting sine waves show a negative-going zero crossing near second-harmonic phase delays of 0° or 360°, so that their maxima and minima lie near the peaks-align or troughs-align waveforms for which the shifts in mean hue are close to their greatest. A sine wave with a negative-going zero crossing at 0° or 360° is of the form $\sin(2\pi ft + \pi)$, where

	L-cone	M-cone
AS		
4 Hz	4.4966	3.3417
6.7 Hz	2.7148	4.5991
10 Hz	3.4829	5.5468
13.3 Hz	4.8218	5.5507
KR		
4 Hz	8.0058	3.2975
6.7 Hz	8.6713	5.6811
10 Hz	4.4239	2.3415
13.3 Hz	2.7826	2.7289

Table 3. Standard errors of sinusoidal fits to the duty-cycle matches for L-cone and M-cone modulations for observers AS and KR at fundamental frequencies from 4 to 13.3 Hz.

f is frequency in Hz and *t* is time in seconds; hence it would be shown in Figure 9 as having a phase of π radians or 180°. The best-fitting phases between 4 and 10 Hz actually lie slightly but consistently below 180°. The best-fitting phase averages about 15° below 180° for AS and 30° below 180° for KR. Thus, the maximum shifts occur at second-harmonic phase delays in the ranges of 60°–75° and 240°–255°, and the minimum shifts occur in the ranges of 330°–345° and 150°–165°. This observation is crucial for discriminating among models of the mechanism underlying the results, and we return to it later.

Experiment IV: Red–green hue discrimination and second-harmonic phase

Introduction

The results of the suprathreshold matches in the previous experiment show a roughly sinusoidal variation of the matching duty cycle with the phase delay of the second harmonic. In this experiment, we relate those suprathreshold-matching data to hue-discrimination data by measuring how the hue shifts from reddish to greenish and from greenish to reddish caused by the second harmonic depend on its phase delay.

Methods

The stimuli again were combinations of first and second harmonics with the second having half the modulation of the first. Observers were presented with pairs of such stimuli in two 5.7° (visual angle) semicircular half-fields separated by 0.6°. The pairs of stimuli differed only in the phase of their second harmonics—that is, their second-harmonic phases were 180° apart (for example, a slowly-on waveform was paired with a slowly-off one, or a peaks-align one with a troughs-align one). The stimuli were cosine-windowed with a time to peak of 1 s.

In the experiment, the modulations of the first and second harmonics were fixed at 10% and 5%, respectively, and the 180° difference in the phases of the second harmonics of paired stimuli was maintained. At each frequency from 6.67 to 13.33 Hz, the observer varied the phase delay of the second harmonic (keeping the 180° phase shift between the second harmonics of the paired stimuli) to find the two narrow ranges of phase delay at which the reddish and greenish hues of the two fields swapped sides, making them indiscriminable at the transition point. Phase steps of 10° or 20° were used. Three runs of five settings each were performed, and the results were averaged. The highest frequencies at which observers could make reliable settings varied with observer and cone type.
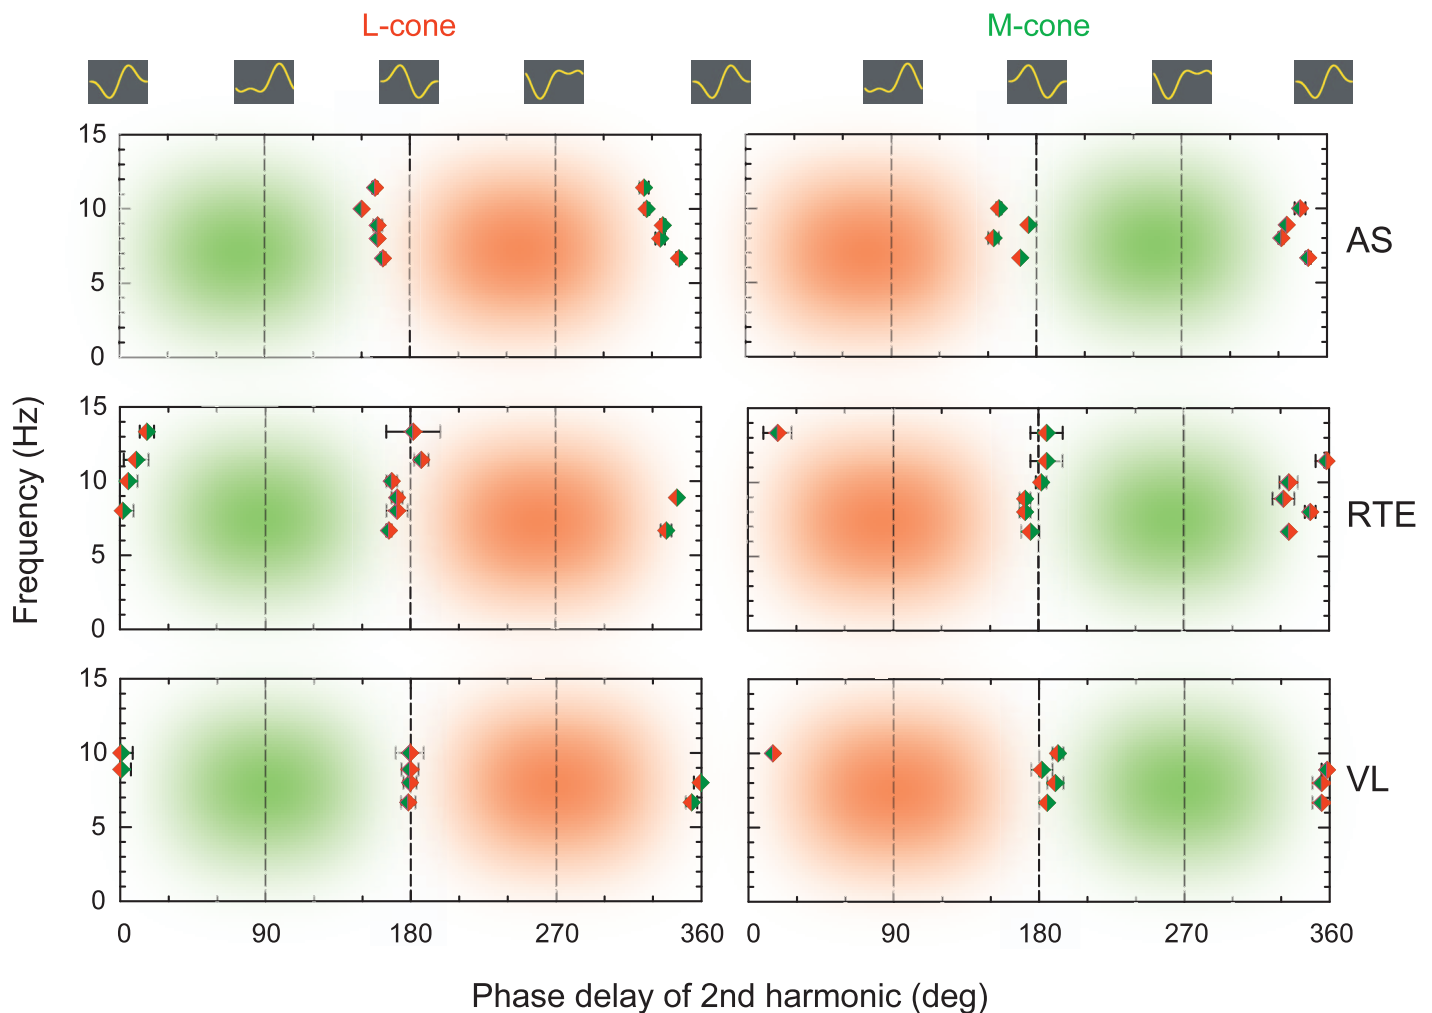
Figure 10. Second-harmonic phase delays plotted along the abscissa for observers AS (top panels), RTE (middle panels), and VL (bottom panels), below and above which the mean hue changes from green to red (green-red diamonds) or from red to green (red-green diamonds) for L-cone (left panels) and M-cone (right panels) stimuli. Each set of data was obtained at the frequencies shown on the ordinate. The second-harmonic phase delays for the slowly-off, peaks-align, slowly-on, and troughs-align waveforms of 0°, 90°, 180°, and 270°, respectively, are highlighted by the icons above the top panels and by the dashed vertical lines in each panel (the phase delay of the slowly-off waveform coincides with the y-axis). The reddish- and greenish-colored backgrounds indicate the predominant mean hue shift within those areas.

Results

Figure 10 shows discrimination data for three observers. The L- and M-cone data are plotted in the left- and right-hand panels, respectively, for AS (upper panels), RTE (middle panels), and VL (lower panels). In each panel, frequency (Hz) is plotted on the ordinate, and the phase delay of the second harmonic (degrees) is plotted on the abscissa. Both axes are linear. The reddish and greenish colored backgrounds in each panel indicate the predominant mean hue that is seen when the stimulus is flickered. The diamonds plot the second-harmonic phase delay at which the mean hues change from a greenish to a reddish hue (green-red diamonds) or from reddish to greenish (red-green diamonds). Measurements were made at fre-

quencies of 4, 6.7, 8, 10, 11.43, and 13.33 Hz; missing data points indicate those frequencies at which the observer was unable to make reliable measurements, usually because the hue shifts were too indistinct. The icons above the top panels are aligned at the second-harmonic phase delays that correspond to slowly-off, peaks-align, slowly-on, and troughs-align waveforms; and the vertical dashed lines in each panel mark second-harmonic phase delays of 90°, 180°, and 270°.

At those frequencies at which measurements could be made, the transitional phase delays from reddish to greenish hues or vice versa were fairly abrupt and thus relatively easy to set. They lie close to the slowly-on and slowly-off waveforms, but often require a slight phase advance or delay of the second harmonic. This phase deviation varies with frequency in some data, particu-
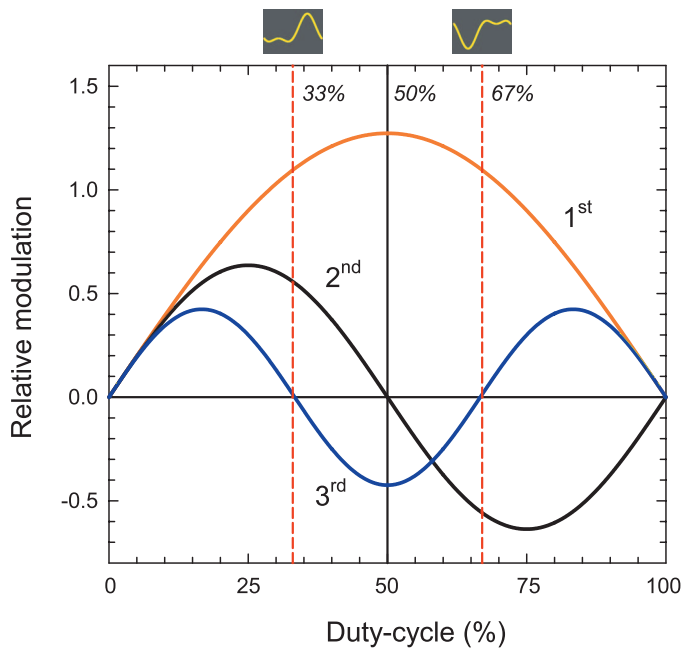
Figure 11. Dependence of the modulation of the first (orange line), second (black line), and third (blue line) harmonics on the duty cycle of a rectangular wave with a constant mean value and unit amplitude. Duty cycle is indicated as a percentage of the period. The vertical solid black line indicates the duty cycle of 50% that corresponds to a square wave. The vertical dashed red lines mark duty cycles of 33% and 67% that have the same ratio of first- to second-harmonic amplitude as the peaks-align and troughs-align two-component waveforms indicated by the icons above the lines.

larly the red-to-green L-cone transitions for AS and RTE, but the shifts are in opposite directions. For AS, the phase transitions occur at second-harmonic phase delays on average about 21° less than the phase delays for slowly-off and slowly-on waveforms (which have second-harmonic phase delays near 0° or 180°). For RTE, the phase transitions are closer to the slowly-off and slowly-on second-harmonic phase delays.

Between the second-harmonic phase delays that form the boundaries of the hue transitions, a mean reddish or greenish hue shift was always apparent, but its salience, especially with suprathreshold modulation, changed with second-harmonic phase delay such that halfway between the boundaries—that is, in the peaks-align and troughs-align conditions—the mean hue was most clear.

Discussion

Given that the initial inspiration for this work and the work reported by Stockman et al. (2017) was the shift in mean hue seen in the slowly changing direction of slowly-off and slowly-on sawtooth waveforms, it was something of a surprise to discover that these

waveforms are far from the optimal ones for producing such hue shifts. In terms of the phase delay of the second harmonic, these hue-discrimination results show that the optimal second-harmonic phase delay differs by 70° to 90° from the slowly-off or slowly-on stimuli. Most other observers (KR and MG, for example) are more like AS in that they consistently see hue shifts for slowly-on and slowly-off sawtooth waveforms in the direction of the slowly changing slope at most frequencies, but for RTE at some frequencies and for VL at all frequencies the slowly-off and slowly-on waveforms are so near to the transitional phase that they do not always produce a clear hue shift in the direction of the slowly changing slope.

These data clearly support the duty-cycle matching experiments described in the previous section, both in terms of the second-harmonic phase delays at which the hue transitions occur and in terms of the transitional phase delays' being roughly independent of frequency. Both sets of data show convincingly that the optimal stimuli for producing the hue shifts are not slowly-off and slowly-on waveforms.

General discussion

Between 0.63 and 2 Hz, observers are able to follow the temporal variations in hue caused by both the first and the second harmonics of the composite waveforms, and can reliably match the reddish and greenish hue excursions produced by different phases of the second harmonic. Consequently, the waveforms that produce the largest or smallest hue excursions correspond to those that produce the largest or smallest peak excursions in cone excitation, which correspond either to the peaks-align waveform (with a second-harmonic phase delay of 90°) or to the troughs-align waveform (with a second-harmonic phase delay of 270°).

At 4 Hz and above, the variations in hue caused by the second harmonic can no longer be followed, and observers report a flickering hue superimposed on a steady mean light, the hue of which varies with the phase of the second harmonic. This change, from being able to follow temporal variations in cone excitation at low frequencies to perceiving chromatic flicker around a shifted mean hue, seems central to the phenomenon.

In the next sections, we focus mainly on the results obtained at and above 4 Hz.

The effect of varying duty cycle

The ability to match the two-component waveforms between 4 and 13.3 Hz against the variable-duty-cycle rectangular wave provided a useful method of investi-
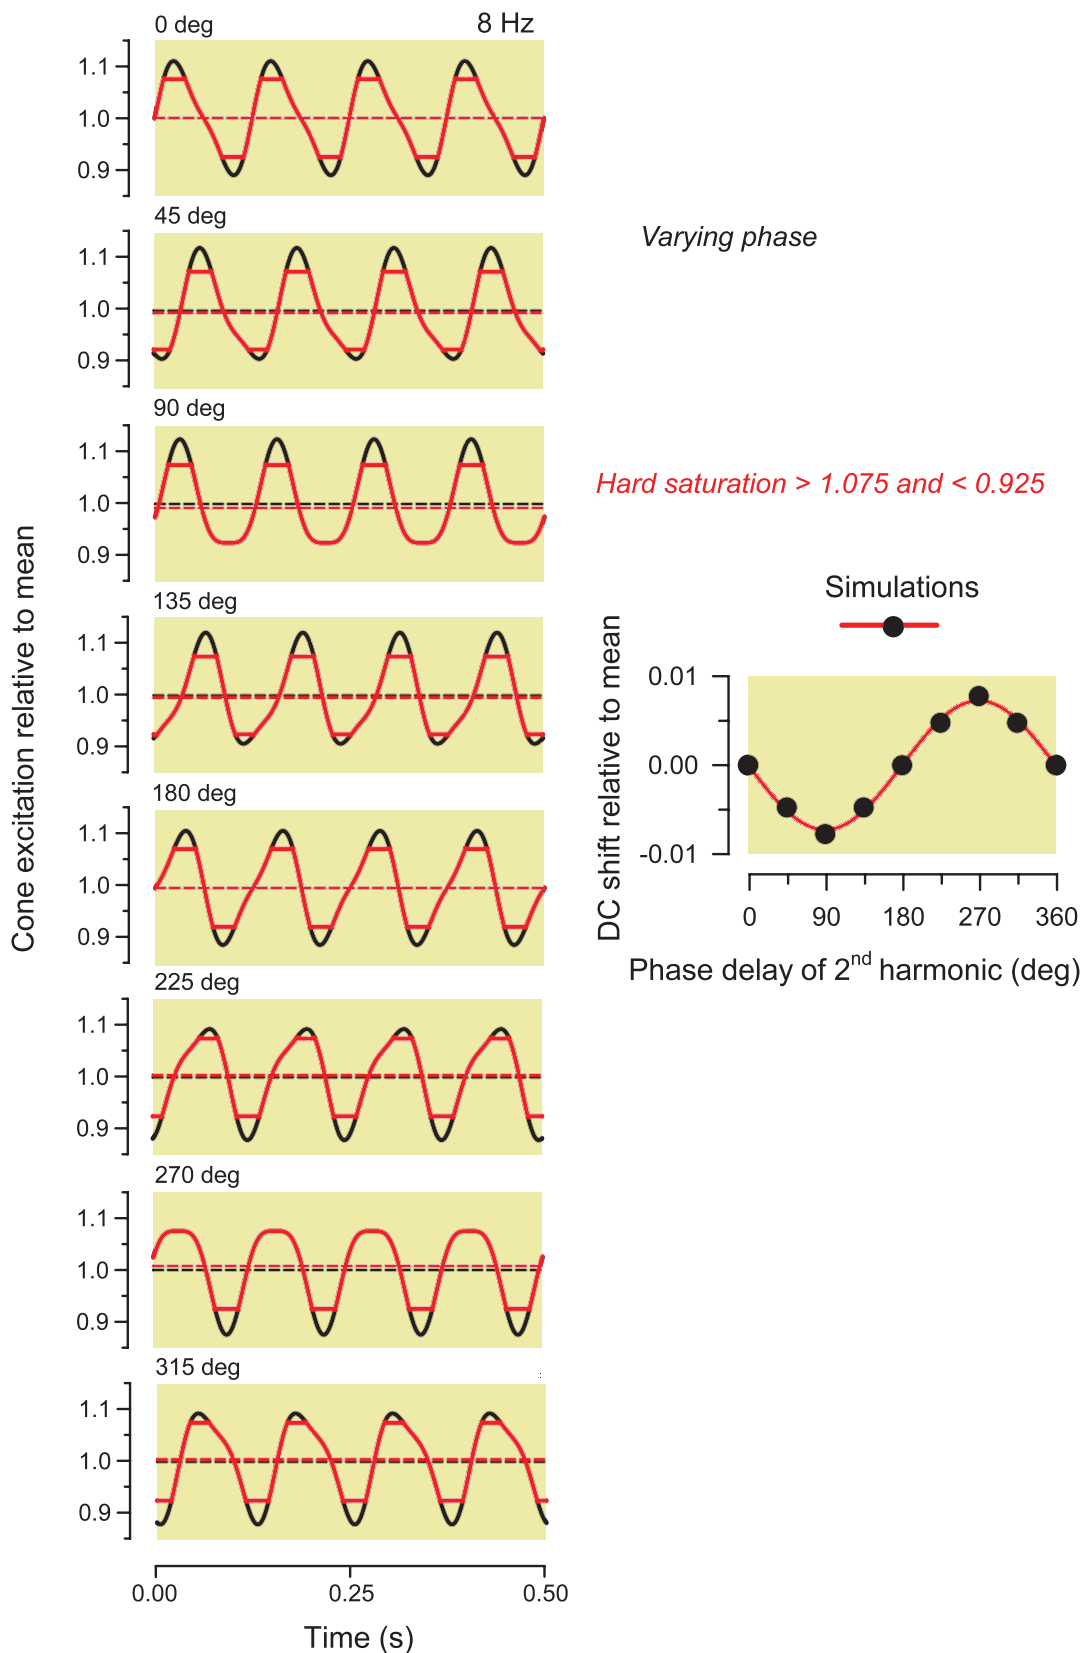
Figure 12. Simulations of the effect of an instantaneous symmetrical saturating nonlinearity on input waveforms with a fundamental frequency of 8 Hz made up of a first harmonic with a modulation of 0.10 and a second harmonic with a modulation of 0.025, and with second-harmonic phase delays that vary in steps of 45° from 0° (top left panel) to 315° (bottom left panel). Each of the eight panels on the left shows the input stimuli as solid black lines and their means as dashed black lines. In the simulations, the input waveforms

→

←

are clipped at ±0.075 above and below the input mean to produce the output waveforms shown by the solid red lines, the means of which are shown by the dashed red lines. The right-hand panel shows the difference between the input and output means as a function of second-harmonic phase delay. The solid circles highlight the simulations shown on the left. The variation in output mean is sinusoidal (solid red line) with a phase delay of the fitted sinusoid relative to the second-harmonic phase delay of **180°**.

gating and quantifying the dependence of the hue appearance on second-harmonic phase delays. As we became more aware of the crucial role of the first and second harmonics, it became apparent that the matches with variable duty cycle had some physical basis.

Figure 11 shows how the relative modulations, shown on the ordinate, of the first (orange line), second (black line), and third (blue line) harmonics of a rectangular waveform of fixed mean and unit modulation vary with duty cycle, shown as a percentage on the abscissa. The duty cycle corresponds to the percentage of the cycle during which the cone excitation is high relative to the mean excitation. For example, for a duty cycle of 50% the variation in cone excitation follows a square wave temporal profile; and for a duty cycle of 75%, the cone excitation is high for 75% of the cycle and low for 25%. Relative modulation, as plotted, is the appropriate variable because as the duty cycle of the stimuli was varied in our experiments, the center point of the excursions was shifted to maintain a constant time-averaged excitation.

A rectangular wave with a duty cycle of 50%—a square wave—has a zero-amplitude second harmonic and thus, neglecting harmonics above the second, is equivalent to a sinusoidal waveform at the fundamental frequency. As can be seen in Figure 11, adjusting the duty cycle away from 50% adds increasing amounts of the second harmonic but in opposite phase above and below the 50% duty cycle. Thus, in terms of their first and second harmonics, the comparison waveforms are either in peaks-align phase (duty cycle less than 50%) or troughs-align phase (duty cycle more than 50%), with duty cycle mainly varying the relative amplitude of the second harmonic. As indicated by the dashed red vertical lines and icons in Figure 11, a rectangular wave with a duty cycle of 33% has the same relative amplitudes (1:0.5) of the first- and second-harmonic components with the same second-harmonic phase delay of 90° as our peaks-align waveform, while one with a duty-cycle of 67% has the same components in the same relative phase (270°) and relative amplitude as a troughs-align waveform. If the nonlinear process that generates the shift in mean hue lies after the temporal filtering in the chromatic pathway, then we need consider mainly only the first and second harmonics of the waveforms, in which case rectangular waveforms with duty cycles of 67% and 33% should match the peaks-align and troughs-align waveforms, since after filtering they are equivalent.

What is surprising, then, is that composite waveforms with second-harmonic phases other than 90° and 270°, which produce dissimilar forms after filtering, are matched satisfactorily against a variable-duty-cycle rectangular wave. We think that the matches are possible because varying the duty cycle but not the mean of the rectangular wave alters the mean hue appearance. In general, increasing the duty cycle above 50% shifts the perceptually associated hue in one direction, while decreasing it below 50% shifts the hue in the opposite direction, so that the hue shifts generated by the composite first- and second-harmonic waveforms can be matched. The second-harmonic phases that produced no mean hue excursions in the combined first- and second-harmonic test waveforms could be matched by square waves of close to 50% duty cycle (see Figures 7 and 8). As expected, those matches occur close to the phase-delay boundaries in Figure 10 where the mean hue changes from reddish to greenish or vice versa.

The general finding that rectangular waves match the two-component waveforms is important because it requires that, in any appropriate model, the two-component waveforms and the rectangular waveforms of the chosen duty cycle produce changes in the mean output in the appropriate hue directions.

Phase dependencies and model predictions

Between 4 and 13.33 Hz, waveforms composed of first- and second-harmonic components are seen as chromatic flicker superimposed on a steady light, the hue of which varies with the phase of the second harmonic (or, in the case of the matching rectangular waveforms, with duty cycle). We have simulated the effects of varying the second-harmonic phase on the output of two explanatory models: a saturating nonlinearity and a slew-rate-limiting nonlinearity. These and other simulations were carried out using MATLAB and Simulink.

Figures 12 and 13 show illustrative simulations of the effects of varying second-harmonic phase on the outputs of a saturating nonlinearity that clips the chromatic signal symmetrically above and below the mean (Figure 12) and a slew-rate nonlinearity that limits the rate of change of the chromatic signal to, arbitrarily, ±2 units/s (Figure 13). In the simulations,
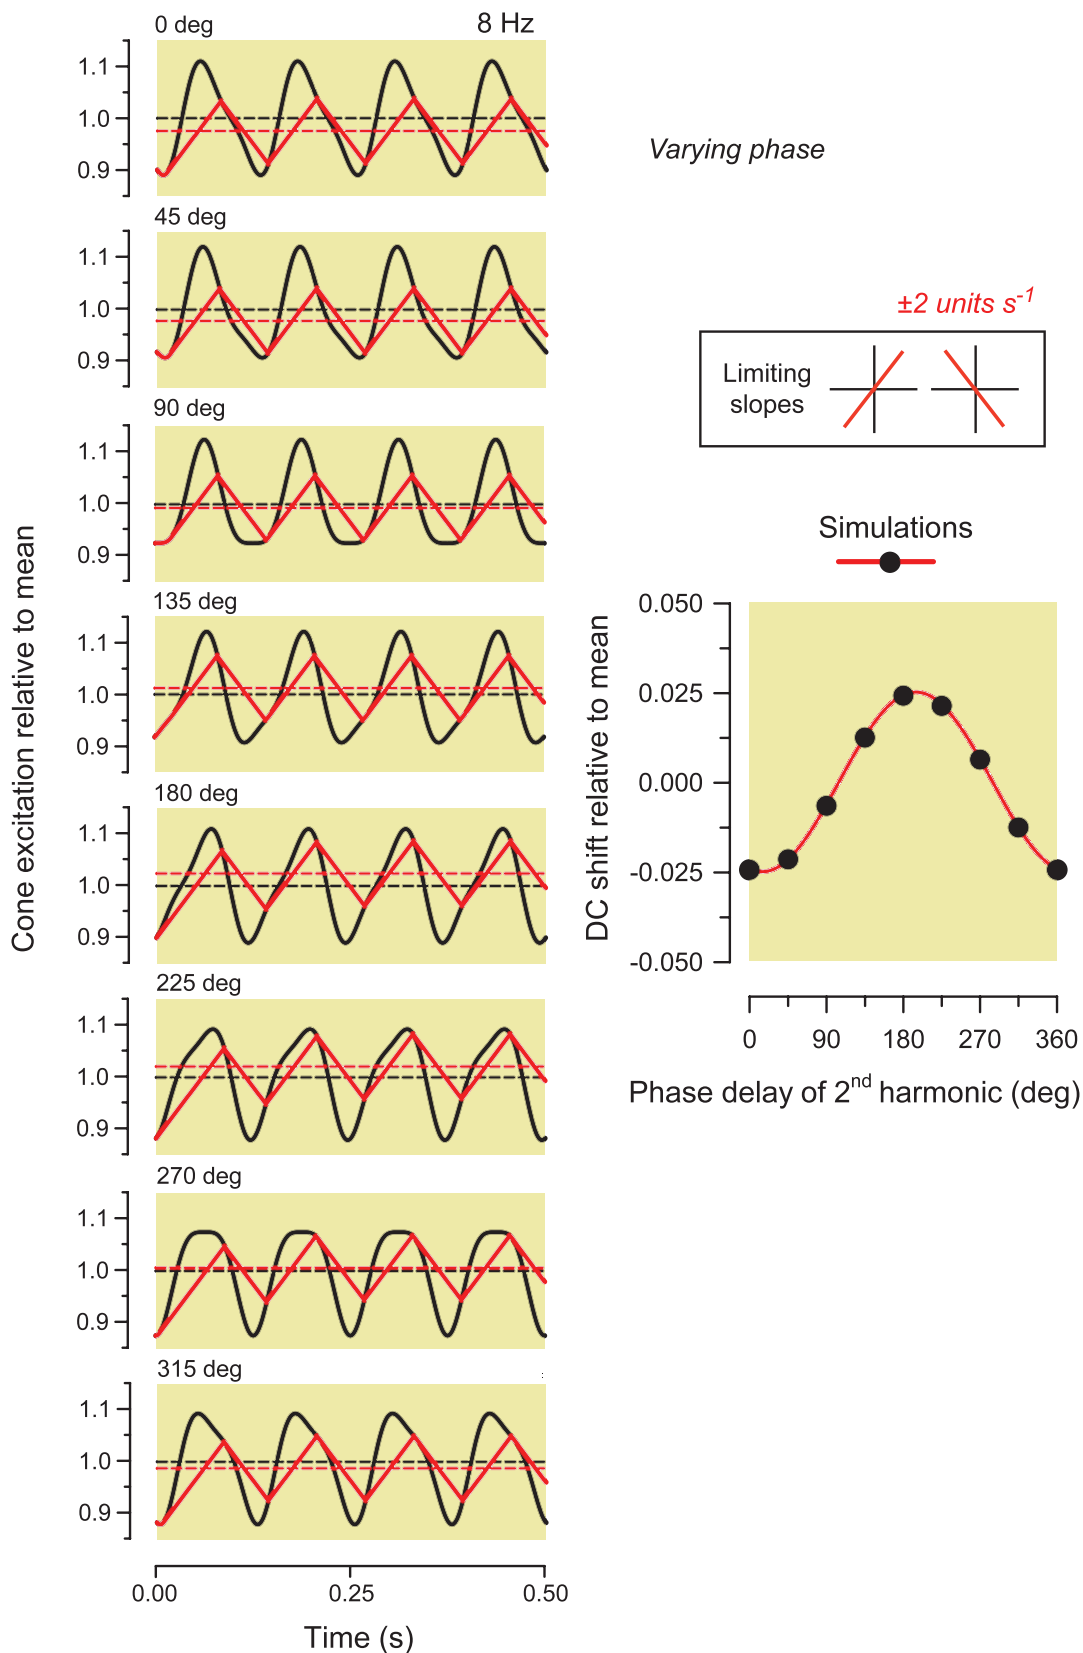
Figure 13. Simulations of the effect of a slew rate of ±2 units/s on the same input waveforms used in the simulations shown in Figure 12. As in Figure 12, the input waveforms and their means are plotted in the left-hand panels as the solid and dashed black lines, respectively. In these simulations, the rate of change of the input waveforms is limited to ±2 units/s to produce the output waveforms shown by the solid red lines, the means of which are shown by the dashed red lines. The right-hand panel shows the

→

←

difference between the input and output means as a function of second-harmonic phase delay. The solid circles highlight the simulations shown on the left. The variation in output mean is sinusoidal (solid red line) with a phase delay of the fitted sinusoid relative to the second-harmonic phase delay of **105**°.

---

the fundamental frequency of the input waveforms is 8 Hz and the first and second harmonics had modulations of 0.10 and 0.025, respectively. (Sawtooth waveforms have second-harmonic modulations that are half that of the first harmonic, but since the chromatic filter affects the sawtooth waveforms at the input to a late nonlinearity, we have reduced the modulation of the second harmonic by an additional 50% to simulate the effects of the late chromatic filter; Stockman et al., 2017.) If, instead, the nonlinearity precedes the filter and the second harmonic is larger, the conclusions do not substantially change; the sizes, but not the directions, of the predicted hue shifts change.

The response of a saturating system to varying second-harmonic phase

The simulations in the left column of Figure 12 show waveforms clipped by a symmetrical hard limit (saturation) set to limit excursions to no more than ±0.075 units from the mean. The input waveforms are shown in black and the output waveforms are shown in red. Their mean levels are shown as the black and red dashed lines, respectively. The waveforms have a fundamental frequency of 8 Hz and are the sum of the first harmonic (with a modulation of 0.1) and the second harmonic (with a modulation of 0.025). The phase delay of the second harmonic ranges in steps of 45° from 0° in the top panel to 315° in the bottom panel. Thus, the first, third, fifth, and seventh panels show the slowly-off, peaks-align, slowly-on, and troughs-align waveforms, respectively.

For the slowly-off (0°, panel 1) and slowly-on (180°, panel 5) waveforms, the excursions from the mean are symmetrical and the symmetrical saturation produces no DC shift and hence predicts no hue shift. For other phase delays the input is asymmetrical, and the saturating nonlinearity alters the mean output signal. As shown in the right-hand panel, the variation of the output mean as a function of the phase delay of second harmonic forms a sinusoid in negative sine phase—that is, a sine wave with a phase of **180**°. Note that to distinguish the phase of the simulated sinusoidal variation from the second-harmonic phase itself we have used **bold** type. The greatest deviations in the mean output occur for the peaks-align (90°, panel 3) and troughs-align (270°, panel 7) waveforms, and the shifts, in a direction opposite from that of the

maximum excursions, are in the direction of our observed hue shifts.

The response of a slew-rate-limited system to varying second-harmonic phase

The simulations in the left column of Figure 13 show the same input waveforms as in Figure 12 but with the rate of change of signal limited to a slew rate of ±2 units/s. Again, the input waveforms are shown in black and the output waveforms are shown in red; their mean levels are shown as the black and red dashed lines, respectively. As in Figure 12, the panels again show the effect of varying the phase of the second harmonic, whose amplitude is again one quarter that of the fundamental. The slew-rate limit (±2 units/s) restricts the steeper slope of each waveform more than the shallower slope and thus shifts the mean output in the direction of the shallower slope when the falling and rising slopes are asymmetrical (see also the slew-rate simulations in our companion article; Stockman et al., 2017). The DC hue shift as a function of the second-harmonic phase is again shown in the right-hand panel. When plotted against the phase delay of the second harmonic, the DC shift is again sinusoidal in form, but for this simulation the sinusoid has a phase of **105**°— that is, it is almost in negative cosine phase. The maximum shifts occur at second-harmonic phase delays of 15° and 195° (near the phases for slowly-off and slowly-on sawtooth waveforms). The minimum shift occurs at 105° and 285°.

We can summarize several general conclusions that can be drawn from the simulations:

1. A saturating nonlinearity that is symmetrical around the mean does not alter the mean output if its input waveforms are slowly-off (0°) or slowly-on (180°), but produces the largest shifts in the mean when the inputs to the nonlinearity are peaks-align (90°) or troughs-align (270°). The DC shifts are in the directions that we observe. The variation in the mean shift with second-harmonic phase is approximately sinusoidal in form and is in negative sine phase—that is, with a phase of **180**° (Figure 12, right-hand panel).

2. A slew-rate limit (of ±2 units/s) that follows slower waveform changes more closely than fast changes produces the maximum shifts in the mean (DC shifts) when the input waveforms have second-harmonic phase delays of 15° (near slowly-

off) and 195° (near slowly-on), and produces the minimum shifts at delays of 105° (near peaks-align) and 285° (near troughs-align). The variation in the mean shift with the second-harmonic phase is again sinusoidal in form with a phase of **105°** (Figure 13, right-hand panel). The phase that produces the maximum shifts will vary slightly with the relative amplitude of the second harmonic but will always lie close to the slowly-off and slowly-on waveforms.

3. The effects of a slew-rate limit can be approximated by introducing a differentiator before a symmetrical saturating nonlinearity. Differentiation increases the amplitude of the second harmonic (relative to the amplitude of the first) and, although it advances the phase of both components by 90°, delays the phase of the second harmonic relative to the first by 90°. From Figure 12, we can predict that the unfiltered input waveforms that produce the greatest mean shifts for a differentiator followed by a saturating nonlinearity will be the slowly-off and slowly-on waveforms. The variation in the mean shift as a function of second-harmonic phase would be sinusoidal with a phase at **90°**—a behavior roughly similar to that of a slew-rate-limiting nonlinearity without prior differentiation.

4. Note that introducing an integrator instead of a differentiator before the nonlinearity will cause a phase advance of the second harmonic by 90° relative to the first harmonic and reduce the amplitude of the second relative to that of the first. If the nonlinearity is a saturating one, however, this advance (from, for example, slowly-off to troughs-align) would predict a mean hue shift in the wrong direction. The direction of hue change would be correct if the nonlinearity were expansive. Our previous work, however, suggests that the late nonlinearity is saturating or compressive (Petrova et al., 2013).

We next turn to our results.

Second-harmonic phase dependencies with respect to the input waveforms

The discrimination data (Figure 10) show that the hue transitions that coincide with the minimum mean hue shifts occur at second-harmonic phase delays of about 159° and 339° for AS, 174° and 354° for RTE, and 182° and 362° for VL. The hue-matching data for AS and KR in Figures 7 and 8 are consistent with the hue-discrimination results for AS, since they show that minimum hue shifts occur near phase delays of 150°–165° and 330°–345°.

On balance, if we average across these estimates (given that the phase delays change very little with frequency), the minimum shifts in hue occur when the second-harmonic phase delays in the stimulus are near 165° and 345°. The maximum shifts in mean hue occur, correspondingly, at second-harmonic phase delays that differ from the minimum-shift ones by approximately 90°—thus, near 75° and 255°. With respect to the input waveforms, these results show that the chromatic pathway as a whole is not strictly a slew-rate-limited system, since for a slew-rate limit the phase delays producing the maximum and minimum hue shifts should be near 0° and 180° (or 15° and 195°), corresponding to slowly-off and slowly-on waveforms, respectively (see Figure 13). With the exception of VL, for most observers such waveforms do change in mean hue in the direction of the slow phase—as predicted by a slew-rate model. However, as can be seen in Figures 7 and 8, the maximum hue shifts occur for input waveforms in which the second-harmonic components are phase delayed by about 75°, which is 75° away from where the slew-rate-limit model predicts they should be.

If we assume that the stimulus conditions at the visual input that cause the maximum hue shifts are also the conditions at the input to the nonlinear mechanism, then the mechanism is much closer to a saturating system than a slew-rate-limiting mechanism, since the second-harmonic phase delays that produce the maximum and minimum hue shifts are only 15° less than the phase delays of 90° and 270° (corresponding to peaks-align and troughs-align waveforms, respectively) that, according to our simulations, produce the maximum hue shifts in a saturating system (see Figure 12). Interestingly, for one observer (VL) the optimal input stimuli are peaks-align and troughs-align waveforms. Thus, for him the system as a whole is a saturating one.

Second-harmonic phase dependencies at the nonlinearity

The conclusions in the previous section considered the second-harmonic phase delays at the input to the visual system. However, if we want to understand the nature of the nonlinear mechanism that causes the mean hue shifts within the chromatic pathway, we need to know how the waveforms at the input to the visual system are changed in amplitude and particularly in phase before they arrive at the hue-shifting nonlinear mechanism. As illustrated by the simulations, knowing the inputs that produce the maximal and minimal hue shifts at the nonlinearity is especially important for discriminating among several possible nonlinear explanations of the hue-shift phenomenon. Thus, it is

clear that we need to know the amplitude and phase characteristics of any filtering that precedes the nonlinear stage where the hue shift is produced. Although this is an ambitious goal, we can draw some general conclusions.

Any filtering must account for the phase delay data that characterize the hue-shift effects shown in Figures 7, 8, and 10. Those data show that the relative first- and second-harmonic phases that produce the maximal and minimal hue shifts are almost independent of frequency between fundamental frequencies of 4 and 13.33 Hz. Thus, irrespective of the type of nonlinearity that causes the hue shift, the phase characteristics of the filter must be such that the relative phase of the first and second harmonics between fundamental frequencies of 4 and 13.33 Hz is roughly constant. This places a severe constraint on any candidate filter, since it suggests that the phase is approximately linear with frequency (i.e., consistent with a fixed time delay) from 4 to 26.67 Hz.

In principle, the filter could be consistent with either the slew-rate or the saturation model. It would be consistent with the slew-rate model if over its fixed range the filter were to advance the relative phase of the second harmonic by about 75° at each frequency. Then the waveforms that produce the maximal hue shifts at the visual input (with second-harmonic phase delays of 75° and 255°) will arrive at the nonlinearity with the second-harmonic phase delays of the slowly-off and slowly-on waveforms (0° and 180°), which are close to the optimal stimuli for the slew-rate-limit model. On the other hand, the filter would be consistent with the saturation model if over its fixed range it were to delay the relative phase of the second harmonic by about 15° at each frequency. Then the waveforms that produce the maximal hue shifts at the visual input will arrive at the nonlinearity with the second-harmonic phase delays of the peaks-align and troughs-align waveforms (90° and 270°), which are the optimal stimuli for the saturation model.

Crucially, if the nonlinearity is late and follows the chromatic filter, then that filter should also be consistent with chromatic TCSF data plotted as orange diamonds in Figure 2 of our companion article for KR and JA (Stockman et al., 2017). Note that a low-pass filter that fits the chromatic TCSF, although it delays the phase of both the first and second harmonics, actually causes a *relative* phase advance of the second harmonic.

We have been able to construct a plausible filter that is consistent with the chromatic TCSFs and causes the 75° phase advance of the second harmonic over the 4- to 13.33-Hz range required for consistency with a central slew-rate mechanism. By contrast, we could not construct a plausible filter that is consistent with the chromatic TCSFs and causes the 15° phase delay of the second harmonic required for consistency with a central saturating nonlinearity. If the nonlinearity is central, we conclude that it is more likely to be a slew-rate-limiting mechanism. Filters that produce intermediate phase changes between a 75° phase advance and the 15° phase delay are also possible, but we find that the optimal fit is close to a 70° phase advance.

We can also construct a band-pass filter that will produce the 15° phase delay of the second harmonic required for consistency with a saturating nonlinearity, but such a filter is inconsistent with the low-pass chromatic TCSFs. Thus, if the nonlinearity is a saturating one, it is likely to be early in the chromatic pathway at a stage where the filtering is still effectively band-pass (see Stockman et al., 2014) prior to some form of late filtering needed to account for the chromatic TCSFs.

Further research will be required to determine the nature of the internal filtering and whether the nonlinearity is early or late.

Corresponding author: Andrew Stockman.
Email: a.stockman@ucl.ac.uk.
Address: University College London Institute of Ophthalmology, University College London, London, United Kingdom.

References

de Lange, H. (1958). Research into the dynamic nature of the human fovea-cortex systems with intermittent and modulated light. II. Phase shift in brightness and delay in color perception. *Journal of the Optical Society of America*, *48*, 784–789.

Estévez, O., & Spekreijse, H. (1982). The "silent substitution" method in visual research. *Vision Research*, *22*, 681–691.

Kelly, D. H., & van Norren, D. (1977). Two-band model of heterochromatic flicker. *Journal of the Optical Society of America*, *67*(8), 1081–1091.

King-Smith, P. E. (1975). Visual detection analysed in terms of luminance and chromatic signals. *Nature*, *255*(5503), 69–70.

King-Smith, P. E., & Carden, D. (1976). Luminance and opponent-color contributions to visual detection and adaptation and to temporal and spatial integration. *Journal of the Optical Society of America*, *66*(7), 709–717.

Kvalseth, T. O. (1985). Cautionary note about R^2. *The American Statistician*, *39*(4), 279–285.

Lee, B. B., Sun, H., & Zucchini, W. (2007). The temporal properties of the response of macaque ganglion cells and central mechanisms of flicker detection. *Journal of Vision*, *7*(14):1, 1–16, doi:10.1167/7.14.1. [PubMed] [Article]

Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, *2*(2), 164–168.

Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, *11*(2), 431–441.

Metha, A. B., & Mullen, K. T. (1996). Temporal mechanisms underlying flicker detection and identification for red-green and achromatic stimuli. *Journal of the Optical Society of America A*, *13*(10), 1969–1980.

Noorlander, C., Heuts, M. J. G., & Koenderink, J. J. (1981). Sensitivity to spatiotemporal combined luminance and chromaticity contrast. *Journal of the Optical Society of America*, *71*(4), 453–459.

Petrova, D., Henning, G. B., & Stockman, A. (2013). The temporal characteristics of the early and late stages of the L- and M-cone pathways that signal color. *Journal of Vision*, *13*(4):2, 1–26, doi:10.1167/13.4.2. [PubMed] [Article]

Regan, D., & Tyler, C. W. (1971). Some dynamic features of colour vision. *Vision Research*, *11*(11), 1307–1324.

Smith, V. C., Bowen, R. W., & Pokorny, J. (1984). Threshold temporal integration of chromatic stimuli. *Vision Research*, *24*(7), 653–660.

Spiess, A.-N., & Neumeyer, N. (2010). An evaluation of R^2 as an inadequate measure for nonlinear models in pharmacological and biochemical research: A Monte Carlo approach. *BMC Pharmacology*, *10*(6), 1–11.

Sternheim, C. E., Stromeyer, C. F., III, & Khoo, M. C. K. (1979). Visibility of chromatic flicker upon spectrally mixed adapting fields. *Vision Research*, *19*(2), 175–183.

Stockman, A., Henning, G. B., West, P., Rider, A. T., Smithson, H. E., & Ripamonti, C. (2017). Hue shifts produced by temporal asymmetries in chromatic signals. *Journal of Vision*, *17*(9):2, 1–20, doi:10.1167/17.9.2.

Stockman, A., MacLeod, D. I. A., & DePriest, D. D. (1991). The temporal properties of the human short-wave photoreceptors and their associated pathways. *Vision Research*, *31*(2), 189–208.

Stockman, A., Petrova, D., & Henning, G. B. (2014). Color and brightness encoded in a common L- and M-cone pathway with expansive and compressive nonlinearities? *Journal of Vision*, *14*(3):1, 1–32, doi:10.1167/14.3.1. [PubMed] [Article]

Stockman, A., & Sharpe, L. T. (2000). Spectral sensitivities of the middle- and long-wavelength sensitive cones derived from measurements in observers of known genotype. *Vision Research*, *40*(13), 1711–1737.

Tolhurst, D. J. (1977). Colour-coding properties of sustained and transient channels in human vision. *Nature*, *266*(5599), 266–268.

Watson, A. B., & Yellott, J. I. (2012). A unified formula for light-adapted pupil size. *Journal of Vision*, *12*(10):12, 1–16, doi:10.1167/12.10.12. [PubMed] [Article]

Yeh, T., Lee, B. B., & Kremers, J. (1995). Temporal response of ganglion cells of the macaque retina to cone-specific modulation. *Journal of the Optical Society America A*, *12*(3), 456–464.