

Chapter 8: Dynamics of perceptual bi-stability: plaids and binocular rivalry compared

Nava Rubin¹ and Jean-Michel Hupé^{1,2}

1- Center for Neural Science, New York University

2- Present address: Centre de Recherche Cerveau et Cognition, CNRS-UPS UMR 5549, Université Paul Sabatier

This chapter presents a set of findings about the dynamics of bi-stable alternations in the perception of plaids. A plaid is a pattern comprised of two superimposed gratings. When the plaid is set in motion, it becomes an ambiguous stimulus: it can be seen as one pattern moving in a single direction (“coherency”), or as two gratings, each moving in different directions with one sliding over the other (“transparency”). The two possibilities are illustrated in Figure 8.1A. In prolonged viewing, the perception of moving plaids switches back and forth between the coherent and the transparent interpretations – a classic example of perceptual bi-stability. Hans Wallach, who was the first to use superimposed gratings to study motion perception, discussed their bi-stability at some length (Wallach, 1935, Wallach, 1976, Wuerger, Shapley & Rubin, 1996). However, in the modern literature this aspect of plaid perception has been virtually forgotten. Since Adelson and Movshon (1982) re-introduced plaids as a tool for studying motion processing, researchers have been using brief presentation 2AFC methods, overlooking the dynamical aspect (but see von Grunau & Dubé, 1993). We therefore decided to give the bi-stable alternations in plaid perception a closer look.

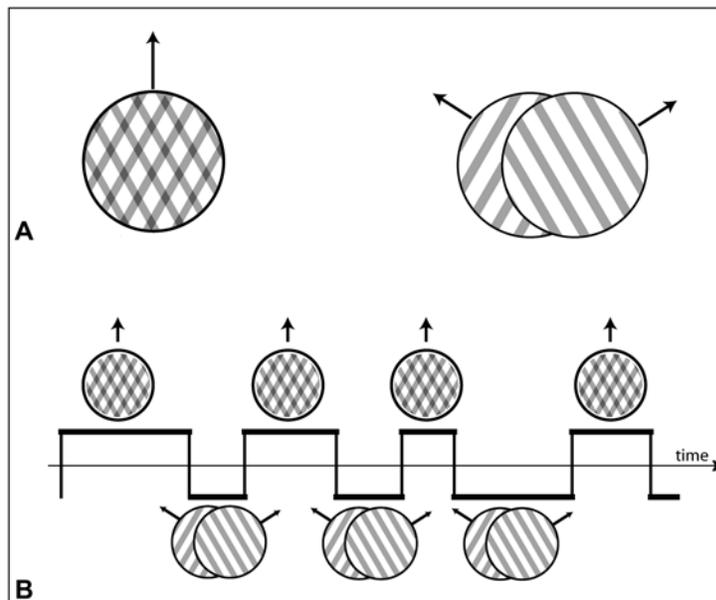


Figure 8.1 A, The two possible interpretations of a moving plaid: a coherently moving pattern (left), or two gratings sliding transparently over each other (right; the lateral offset between the gratings is for illustrative purposes, to show the impression of a slight depth difference between the gratings that accompanies the percept of transparent motion).

B, The dynamics of bi-stable alternations in the perception of moving plaids can be measured by requesting observers to continually indicate their percept (analogous to the paradigm commonly used in binocular rivalry).

We had two complementary goals in this research. The first was to develop methods for studying the dynamics of perceptual alternations in plaids in order to further our understanding of global motion computation. Like all visual information, motion information is fragmented in the first stages of cortical processing, distributed over many neurons with small spatial receptive fields. Therefore, the brain must combine ‘local motion cues’ into ‘global motion percepts’. Global motion computation is non-trivial because real world scenes contain multiple, often overlapping objects that can move in different directions, leading to a complex array of local motion measurements. Thus, on the one hand, there is a need to *integrate* local motion signals that arise from the same object, while on the other hand it is necessary to *segment* motion cues that arise from different objects (Braddick, 1993). Plaids provide a clear illustration of those conflicting demands: in the “coherent” interpretation the integration process is dominant, while in the “transparent” interpretation the motion segmentation process is stronger (the grating components of the plaid are segmented from each other). However, most studies of plaids have used methods involving brief-presentations. By contrast, in other domains – most notably binocular rivalry – the dynamics of bi-stable alternations associated with prolonged viewing have revealed important insights about the underlying mechanisms. For example, the periods spent perceiving each percept vary systematically when the relative strength of the stimuli is manipulated (e.g., by changing the contrast of one or both monocular stimuli; cf. Chapter 1, this volume). We hypothesized that in developing dynamics-based methods (akin to those used in binocular rivalry) to measure the likelihood of the coherent and transparent percepts in plaids, we might shed new light on the mechanisms underlying motion integration, segmentation, and the interplay between them. Indeed, as we shall see below, our results indicate that dynamics-based measures can be more sensitive than brief-presentation measures, thereby revealing effects which were not known until now.

As research towards the first goal progressed, a second, complementary goal emerged. If our first goal can be summarized as “using dynamics of bi-stability to learn about plaids”, the second goal can be stated as “using plaids to further our understanding of bi-stability”. At the heart of this second goal is the idea that bi-stability reveals general principles about brain architecture. Consider this fundamental characteristic of bi-stability: when a stimulus has more than one plausible interpretation, observers will alternate between perceiving one interpretation or another, *but they will always experience only one percept at a time*¹. This is such a basic observation, valid in such varied domains of bi-stability (e.g., Necker cubes, binocular rivalry, vase/face illusion), that it is rarely questioned. But why should it be so? Why should we not be able to perceive *both* stimulus interpretations simultaneously?

One answer is that bi-stability reflects real-world constraints: a location in space cannot be occupied by two different objects simultaneously (binocular rivalry); an object cannot be convex and concave at the same time (the Necker cube), and so on. But this can only be one part of the answer -- the other part has to do with brain architecture. The very fact of bi-stability suggests that the brain has built-in mechanisms to enforce *mutual exclusivity*: given a stimulus with more than one plausible interpretation, the neural representation of only one of those interpretations is allowed to dominate at each moment². It is unlikely that the brain developed specialized mechanisms for bi-stability just so that it can deal with the rare cases of deeply ambiguous stimuli (which are typically encountered only in the lab). Probably, bi-stability – and its implied principle of mutual exclusivity – occurs as a result of brain architecture which

evolved to deal with the far more common situations of ‘weak ambiguity’ present in many sensory stimuli. Normally, a wealth of cues in the environment render one interpretation much more likely than others – but the fact that competing interpretations are seldom experienced (outside the lab) is nevertheless not obvious a priori.

How the brain achieves this uniqueness of perception requires explanation, and bi-stability may offer a window to the underlying mechanisms. As we observed, there are some commonalities between the dynamics of bi-stability in plaids and those reported in other domains (mainly binocular rivalry). These commonalities led us to hypothesize the existence of general principles governing how the brain implements mutual exclusivity. In more concrete terms, then, our second goal is to study the commonalities – as well as differences – of bi-stability in plaids and other domains, how they constrain models of the underlying mechanisms.

8.1 Constancy over time of the average duration of bi-stable alternations

Do the average durations of coherency and transparency epochs during plaid perception show any consistent trends over time? There are two reasons to ask this question. First, there have been many references to “adaptation” or “satiation” in the literature as playing a role in plaid perception (e.g., Adelson & Movshon, 1982, Wallach, 1935). One might therefore expect to find some fingerprint of adaptation in the dynamics of alternations – for example, a slowing down of the alternations in prolonged observations. (Weakening the competing stimuli leads to slower alternations in binocular rivalry; if adaptation has the effect of weakening the perceived interpretation, a consequent slowing down of the alternations might be expected.) The second reason to examine how the average durations behave is methodological: if they are stable over time, this would facilitate deriving dynamics-based measures of the strength of the coherency and transparency percepts. We therefore examined the durations spent perceiving coherency and transparency over very long observation times. Observers watched a moving plaid for five minutes and reported what they perceived (“coherence” or “transparency”) continually by pressing one of two mouse buttons (see Figure 8.1B; cf Hupé & Rubin, 2003 for details of experimental procedures). This procedure was repeated ten times with the same stimulus, but with very long breaks between consecutive trials: there were at most two trials per day (one in the morning and one in the evening).

Figure 8.2 shows scatter-plots of epoch durations, separated into coherency (left) and transparency (right) epochs. While there is wide variability between the durations of the epochs in each trial, when all trials are superimposed a clear picture emerges: the *average* durations of the coherent and transparent epochs are very stable over time³. (For more data supporting this observation see Hupé & Rubin, 2003). There is one exception to this statement: the average duration of the very first epoch, which was always coherency, was longer than that of subsequent coherency epochs. This singularity of the first epoch may be why researchers had a subjective impression that the epochs shorten over time (e.g., Wallach, 1935; J.A. Movshon, personal communication), and why they conjectured about the role of adaptation (or “satiation”; Wallach, 1976) in the alternations. However, our data indicate that this subjective impression is misleading, created solely by the first singular epoch. When the first epoch is excluded, the best-fit linear trends show a zero slope for both coherency and transparency. This means that we can

really talk of a “steady state” phase of bi-stable alternations in plaids, which sets in immediately after the first epoch⁴ and consists of epochs which, although variable in duration, are drawn from distributions with stationary means and variances.

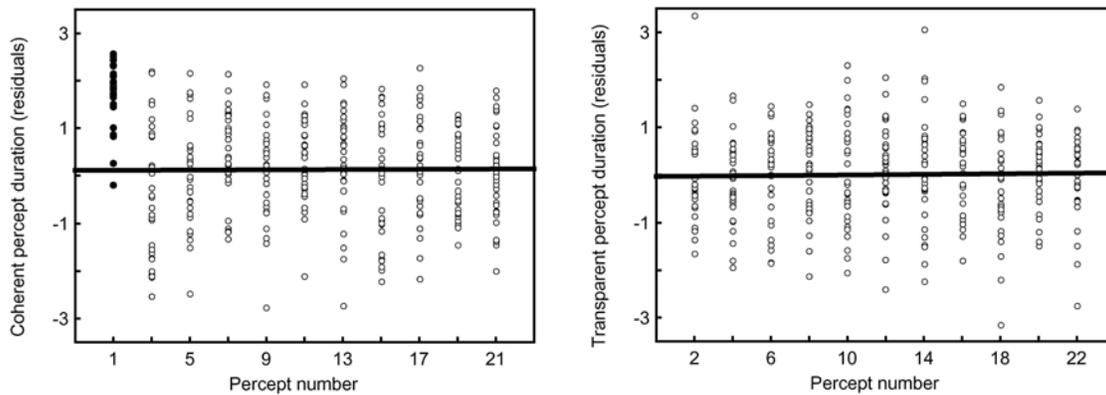


Figure 8.2 Duration of “coherency” (left) and “transparency” (right) epochs in bi-stable alternations of a single plaid stimulus viewed for multiple trials of 5-minutes each (3 observers; data present residual variance after removing inter-subject variability.). The durations are plotted as a function of their ordinal position within each trial. Best-fitting linear trends indicate no significant drift up or down, i.e., the average durations are constant over time (excluding the first coherent epoch, see text). For stimulus parameters and sample scatter-plots of individual observers see Hupé and Rubin (2003).

The stability of the average durations over time has both methodological and theoretical implications. Methodologically, it means that we can derive measures of the relative strength of the coherent and transparent percepts from dynamics data, without concern that those measures may change due to arbitrary factors such as observation time. Specifically, if we denote the average durations spent reporting coherency and transparency by C and T , respectively, then $C/[C+T]$ is the steady-state probability of perceiving the coherent percept, which gives a measure of its relative strength. (Note: the first coherency percept is excluded from C ; it will be treated separately, see below, section 8.3.) This measure is analogous to that used in binocular rivalry studies, the relative time spent perceiving one of the monocular stimuli ($R/[R+L]$), which is known to vary with stimulus strength (e.g., by changing contrast; Levelt, 1968). Indeed, in the next section we show that $C/[C+T]$ varies systematically with manipulations of plaid parameters. Moreover, the dynamics-based measure is sensitive to parametric variations even in regimes where brief-presentation methods suffer from “ceiling” and “floor” effects.

From a theoretical point of view, the stability of the average durations imposes significant constraints on models of bi-stability. Many models assume some form of adaptation of the dominant percept as the main factor that leads to alternations. If such adaptation indeed plays a role, our data suggest that it has to be short lived and its effects cannot accumulate over time.

Otherwise, the average durations would change over time (as they change with stimulus strength). Given the clear results we obtained in plaids, it is interesting to pose the same question also for binocular rivalry stimuli – do the average durations of bi-stable alternations change over time? Previously, there had been reports of lengthening of the durations over time. Lehky (1988) proposed this may be an indirect result of adaptation to contrast (decreased effective contrast leads to weaker stimuli which lead to longer bi-stable durations). If Lehky’s conjecture is correct, then binocular rivalry stimuli that are not affected by contrast should show stable average durations, as we found for plaids. This is not a hypothetical statement: although contrast often affects the rate of rivalry alternations (cf. Chapter 1, this volume), for some stimuli contrast is not a major factor. To test this directly, we used a stimulus introduced by Bossink, Stalmeier & De Weert (1993), consisting of fields of random dots moving in orthogonal directions in the two eyes. For this stimulus, speed has a major effect on the monocular stimulus strength, whereas contrast has only a minor effect. Figure 8.3 shows scatter-plots of the dominance durations for the left-eye and right-eye stimuli for ten repetitions of 5-minute trials (each trial was well separated in time from the others, as before). The results suggest that when extraneous factors (such as contrast-adaptation) are eliminated, the average durations of perceptual alternations in binocular rivalry, as for plaids, are stable over time. (Similar results were obtained from two naive observers; data not shown.) This in turn suggests that, whatever role adaptation plays in causing perceptual alternations, the mechanisms may be similar in different bi-stability domains.

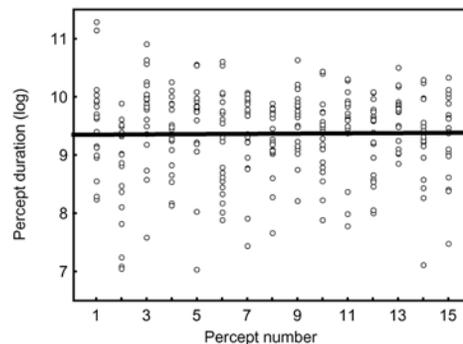


Figure 8.3 Durations of monocular epochs in bi-stable alternations of a binocular rivalry stimulus viewed for multiple trials of 5-minute each (2 observers). The stimuli consisted of random dot surfaces moving in orthogonal directions in the two eyes (cf. Bossink et al. 1993; epochs computed in ‘compound’ method, cf. Mueller and Blake 1989). The average durations of perceptual alternations are constant over time. In contrast to bi-stability in plaids, there is no tendency for the first epoch in binocular rivalry to be longer.

Another way to ask about possible adaptation effects is to check whether there is a correlation between the durations of a given coherency epoch and the subsequent transparency epoch, and between a given coherency epoch and the next coherency epoch (and similarly for a given transparency epoch). Figure 8.4 presents scatter plots for the four pairs. Denoting the *i*-th coherency epoch by C_i and the following transparency epoch by T_i , the data for $\{T_i \rightarrow C_{i+1}\}$ and

$\{T_i \rightarrow T_{i+1}\}$ show no correlation at all, while $\{C_i \rightarrow T_i\}$ and $\{C_i \rightarrow C_{i+1}\}$ show significant but weak correlations (slight tendency for longer C_i 's to be followed by longer T_i 's and then by shorter C_{i+1} 's). Overall, the data suggest that the length of one epoch has little effect on the following epochs. Lehky (1995) performed a similar analysis for binocular rivalry data and found no correlations between the lengths of successive epochs. This suggests that, if adaptation is involved in bi-stability, its influence does not carry over from one epoch to another, since if it did one should observe systematic trends, e.g., long T_i epochs followed by longer than average C_{i+1} epochs and shorter than average T_{i+1} epochs (due to the increased accumulated adaptation of T in epoch i).

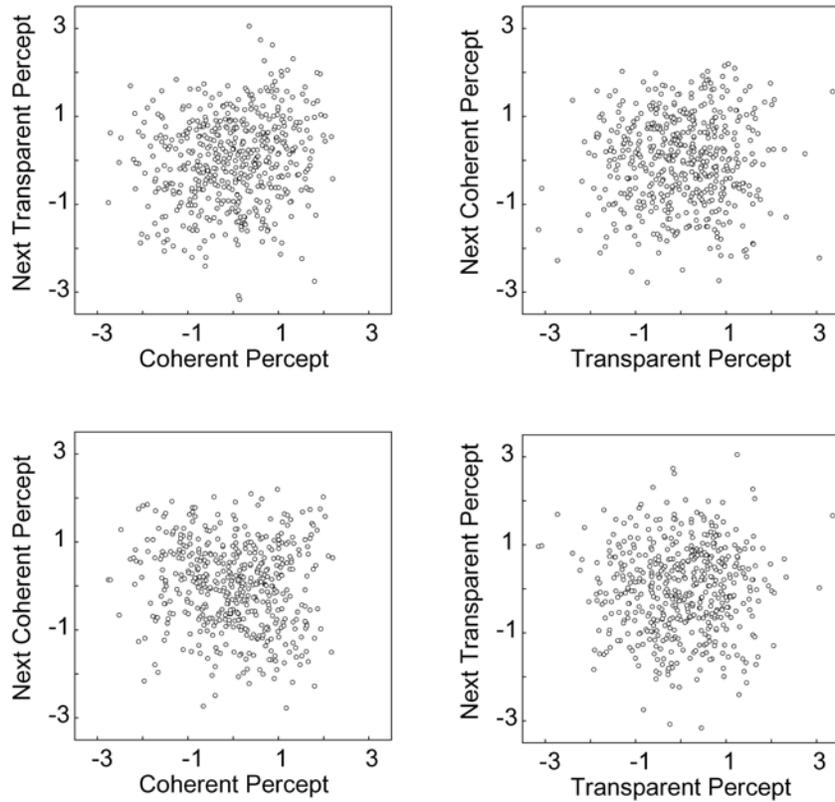


Figure 8.4

Scatter plots of the durations of successive epochs (top, between-type; bottom, within-type). The length of one epoch has little or no effect on the length of the following epochs (see text).

8.2 The distribution of bi-stability durations

Next, we examined the distributions of coherency and transparency durations, and the relations between them. Figure 8.5 shows a histogram of the coherency durations as a function of duration length. As observed in other domains of bi-stability, the distribution resembles a skewed Gaussian, peaking at intermediate values and falling off slowly, with a long ‘tail’. In

other domains of bi-stability (most notably binocular rivalry), Gamma distributions have typically been used to fit the histograms of bi-stable durations (e.g., Leopold & Logothetis, 1996, Levelt, 1968, Logothetis, Leopold & Sheinberg, 1996). For our plaid data, we find that log-normal distributions provide as good a fit, and often better. For the data in Figure 8.5, for example, the best fit log-normal yielded $p < 0.07$ whereas the best fit Gamma yielded $p < 0.2$. Interestingly, Lehky (1995) also reported that for his binocular rivalry data, when compared directly (for the same data set) log-normal distributions provide as good or better fit than Gamma distributions.

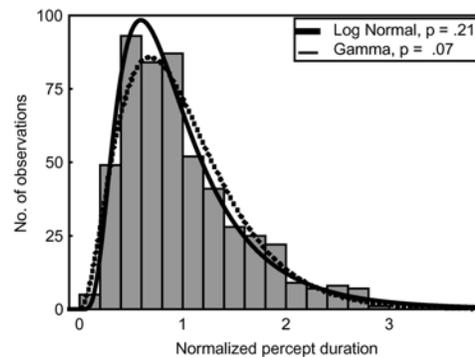


Figure 8.5

Histogram of the durations of coherent periods is better fit by a log-normal than a Gamma distribution. Data from repeated 5-min trials with a single stimulus, three observers (same as in Fig. 2). Data from individual observers were normalized to each observer's mean before pooling.

The idea that the log-normal function provides a better fit to bi-stable alternations data is intriguing, because it suggests that the logarithm of durations of each percept are normally distributed, which has consequences for what might be the underlying mechanisms of the perceptual switches. It is, however, not trivial to obtain data that would clearly and strongly point to one functional form over another. Pooling data across observers, or even just across trials with different stimulus parameters, leads to mixing of values that are drawn from different distributions (i.e., distributions that have different parameters, though the same functional form). To address this problem, researchers often normalize data to observers' individual means before pooling (as we did in Figure 8.5). But as long as the underlying distributions generating the data are not well understood, it is also not clear what effect this procedure has on the subsequent ability to make strong statements about what functional form best fits the data. Furthermore, while normalizing data per observer is common, in most studies data from trials with different stimulus parameters are pooled with no adjustment procedures. We addressed this issue in analyzing data from our experiments with plaids. We collected dynamics data for a large stimulus set, systematically varying plaid parameters that affect the tendency for coherence, and then performed an ANOVA with all parametric manipulations (as well as observer identity) taken into account. The analysis revealed two important points, first, that transforming the raw data to log values led to models with excellent fit to the parametric effects, and second, that the residual variances distributed normally (Hupé & Rubin, 2003). This latter point provides strong support for the claim that the log of durations is indeed normally distributed. In turn, this

supports the use of log-normal distributions to fit histograms of perceptual durations. Theoretical work is needed to assess the implications of these findings for models of bi-stability.

8.3 Dynamics-based measures of the strength of coherency

In a series of experiments, we used the dynamics-based measure $C/[C+T]$ to test the effect of manipulating various plaid parameters on the relative strength of coherency. (Recall that this measure gives the relative time spent perceiving coherency out of the entire observation time.) The angle between the gratings' directions of motion, Alpha, has the most dramatic effect (Fig. 8.6). As Alpha is changed from small to large, perception shifts from being dominated by coherency to being dominated by transparency.

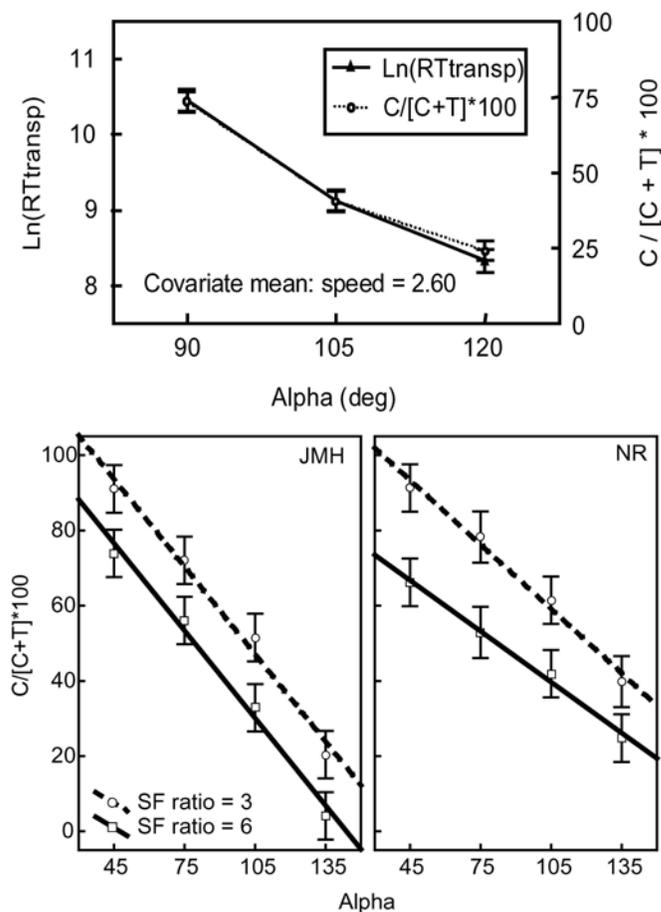


Figure 8.6 The dynamics-based measure $C/[C+T]$ shows a linear relationship to Alpha (the angle between gratings' directions of motion) in a wide range of plaid parameter space. Top panel, rectangular-wave plaids, Four observers. Bottom panels, sine-wave plaids, data from individual observers. For stimulus parameters see Hupé and Rubin (2003).

The importance of the angular separation between motion signals as a strong cue for motion segmentation has a clear ecological basis, since different objects tend to move in independent directions. The effect of Alpha on the probability of perceiving coherency was observed previously, with brief-presentations methods. However, those methods yielded sigmoid-shape curves with rapid transitions between 100% responses “coherency” and 100% responses “transparency”, flanked by wide ranges of Alpha where the responses were constant at those extreme values (e.g., Kim & Wilson, 1993). In contrast, the dynamics-based measure $C/[C+T]$ reveals a gradual, near-linear transition from 0% to 100%, indicating a gradual – not abrupt – change in relative strength between coherency and transparency. In a detailed analysis elsewhere we showed that the sigmoid curves found with brief-presentation methods resulted from “ceiling” and “floor” artifacts inherent to this paradigm (Hupé & Rubin, 2003). Essentially, the brief trials used there (~1 sec) reflect only what happens in the initial observation periods, and those tend to be dominated repeatedly by the same percept (either coherency or transparency), leading to stimuli classified categorically as one or the other. However, when more observation time is allowed (~1 min), bi-stability is observed even when one of the two percepts is strongly dominant in short observations. Thus, the dynamics approach reveals the true underlying effect of increasing Alpha, which is to decrease the relative strength of coherency in a gradual, near-linear manner⁵.

In section 8.1 we noted that for plaids composed of rectangular-wave gratings the very first epoch was always coherency, and its mean duration was longer than that of subsequent coherency epochs. Interestingly, a long initial epoch of coherency was observed even for stimuli which were subsequently dominated by the transparency percept, i.e., for which $C/[C+T]$ was less than 0.5 (e.g., for large Alpha values). This asymmetry between the two competing percepts, coherency and transparency, may be unique plaids⁴. Preliminary findings suggest the asymmetry may be further unique to stimuli where there is ambiguity about the relative depth of the two constituent gratings, so that the transparency ‘percept’ is in fact comprised of two possible percepts which alternate over time (first one of the gratings is seen in front, and then the other). It could be that this further “splitting” of one of the competing percepts (transparency) is what breaks the symmetry between it and coherency, but more research is needed to test this hypothesis further.

Whatever the reason for the prolonged first coherency percept in rectangular-wave plaids, we have been able to use it as an additional measure of the relative strength of coherency. The mean duration of the first epoch can be extracted from the same experiments where observers continually indicate their percept in 1-2 min trials (for measuring $C/[C+T]$; recall that the first coherent epoch is excluded from the calculation of $C/[C+T]$ and therefore the two measures are methodologically independent). It can also be collected more efficiently, in a paradigm where observers are asked to press a button as soon as they see the plaid separate into two transparent gratings, terminating the trial. We therefore termed this measure RT_{transp} , (‘response time to report transparency’; note, however, that this is not a speeded-reaction task: observers were asked not to “try” to see more of one or the other percept, i.e., “passive” viewing instructions). We find that in both cases, RT_{transp} shows the same gradual, linear decrease with Alpha, in agreement with the gradual decrease in the relative strength of coherency revealed by $C/[C+T]$ (Hupé & Rubin, 2003).

The gradual effect of Alpha on the strength of coherency has important implications for models of motion integration and segmentation, since it suggests that there is not a "critical" value of Alpha where the system switches from one interpretation to the other. Several existing models of plaid perception implement a "decision" between coherency and transparency with a switch at such a critical value, mirroring the rapid transition from all-coherency to all-transparency found with brief-presentations (see, e.g., Wilson & Kim, 1994). However, our dynamics data indicate that the underlying transition is actually gradual. This behavior resembles what is found in binocular rivalry, of a gradual change in $R/[R+L]$ as the contrast of one of the monocular stimuli is changed. This type of behavior is explained well by models of binocular rivalry, which assume a continual competition between the two rivaling stimuli, with mutual inhibition insuring that only one is allowed to prevail at any given moment (Blake, 1989, Laing & Chow, 2002, Lehky, 1988). This type of model can be adapted to motion segmentation and integration, and they would naturally give rise to the gradual changes in the relative strength of coherency and transparency in plaids that we observed. We elaborate on this in the next section.

8.4 Levelt's second proposition in binocular rivalry and in plaid perception

No comparison to the dynamics of perceptual alternations in binocular rivalry would be complete without reference to Levelt's second proposition (Levelt, 1968). We have already mentioned that changing the strength of one of the monocular stimuli (e.g., by changing its contrast) affects the relative time spent perceiving that stimulus (see also Chapter 1, this volume). The change is in the direction one would expect intuitively: when the stimulus is strengthened (while the other is kept unchanged), the relative time spent perceiving it increases. Levelt (1968) examined more closely how this increase took place, by studying what happened to the *absolute* durations of epochs perceiving each of the two monocular stimuli. (Recall that while the lengths of the bi-stable epochs vary stochastically from one alternation to the other, their means are stable over time, for a given stimulus. One may therefore study the effect of stimulus manipulations on the length of those mean durations.) Intuitively one might expect that increasing the strength of one stimulus would lead to an increase in its own mean dominance duration, and possibly that there would also be a decrease in the mean dominance duration of the other (competing) stimulus. But when Levelt (1968) tested this he found, surprisingly, that the changes in absolute durations were restricted to those of the other stimulus (the one whose strength was not changed). He summarized this result in what has become known as Levelt's second proposition: "Increase of the stimulus strength in one eye will not affect t^- [the mean dominance durations] for the same eye ... t^-_r [the mean dominance durations of the right-eye stimulus] can only be affected by λ_l [the left-eye stimulus strength], not λ_r ".⁶

Why is Levelt's second proposition surprising? What implicit assumptions does this finding challenge? To answer this, let us consider the simplest model of how the perceptual switches might arise. Assume that the dominant (perceived) stimulus undergoes gradual adaptation, i.e., some form of weakening of the neural activity representing it. At some point the competing percept would become (relatively) stronger and consequently take over. In this simple model, a strengthening of stimulus A would be expected to increase its own dominance duration – because it would take longer for the adaptation to bring it to a level low enough for B (the competing stimulus) to take over. Moreover, this simple model provides no mechanism for the

strength of stimulus A to affect the mean dominance durations of B at all – let alone affect primarily those durations.

There is wide agreement about the crucial ingredients of a model that would account for Levelt's second proposition. Such a model must allow some form of *coupling* between the neural representations of the two percepts. Indeed, most modern models of binocular rivalry use an architecture of reciprocal inhibition between two neuronal populations that represent the rivaling percepts (Blake, 1989, Laing & Chow, 2002, Lehky, 1988). As a consequence, increasing the strength of one stimulus also increases the inhibition it exerts on the competing stimulus, and it is via this inhibitory coupling that one stimulus affects the dominance durations of the competing stimulus.

Levelt's second proposition had significant impact on how binocular rivalry is understood, because it suggested that the bi-stability results from an active process of continual competition between the two monocular stimuli. But the view of bi-stability as an active process of competition remained largely confined to binocular rivalry. In other domains (e.g., ambiguous figures) the view that perceptual alternations result from passive adaptation ('fatigue') of the dominant percept is still common. This is also the case for plaids, where models generally implement a "decision" between the two possible interpretations. The concept of a continual, active competition between the two possible interpretations is virtually absent from the literature.

But the parallels we observed between the dynamics of perceptual alternations in binocular rivalry and in plaids led us to hypothesize that the bi-stable alternations in plaid perception may also arise from an active competition, between the coherency and transparency interpretations. Drawing on the analysis of what Levelt's second proposition implied for binocular rivalry, we reasoned that if one observed a similar phenomenon for plaids, it would be strong evidence for our hypothesis. We therefore asked: if we could change the strength of only one (say the transparent) interpretation, without affecting the other (coherent) interpretation, what would be the effect of this manipulation on the mean dominance durations of the two interpretations?

In binocular rivalry, changing the strength of one of the competing stimuli without affecting the other is most often done by changing the contrast of one monocular stimulus. But in the case of plaids, it is not obvious how to affect only the transparent (or coherent) interpretation. Since the two stimulus interpretations are mediated by the same external image, manipulation of a parameter like contrast would obviously affect both the transparent and the coherent percept. The present lack of thorough understanding of the mechanisms underlying motion integration and segmentation further complicates the task of establishing that we affected only one interpretation. Nevertheless, we were able to find a manipulation which quite certainly changed the strength of only the transparent percept, without affecting the coherent percept. Without going into the details here, it involved switching the depth relationship between the two constituent gratings (by varying the intersections' luminance, not stereoscopic disparity; Hupé and Rubin, manuscript in preparation). When we checked the effect, we found that the average dominance duration of transparency did not change. Instead, when the transparent interpretation

was strengthened, coherency durations decreased markedly. This result is, in effect, a generalization of Levelt's (1968) second proposition to the domain of plaid perception.

The finding that an analogue of Levelt's second proposition can be observed for plaids supports our hypothesis that in the domain of motion perception, too, bi-stable alternations are caused by active competition between the two interpretations – here, the coherent and transparent percepts. Transferring the ideas prevalent in binocular rivalry to motion integration and segmentation, this would suggest an architecture where the neural representations of the coherent and transparent interpretations mutually inhibit each other in a struggle for perceptual dominance. These ideas represent a significant departure from the view common at present, which assumes that integration requires global processing while segmentation is obtained directly via the responses of local motion detectors (Adelson & Movshon, 1982, Movshon, Adelson, Gizzi & Newsome, 1985). Models that implement this approach typically involve a feed-forward stage of “decision” whether or not to integrate the local cues (e.g., Wilson & Kim, 1994). In such models, a need for competition does not arise naturally.

There is, however, evidence that motion segmentation also requires global processing (cf. Braddick, 1993). Indeed, some authors have postulated independence of integration and segmentation mechanisms – e.g., Liden & Pack (1999) proposed to “subdivide the second stage into two parallel computations, one for integration and the other for segmentation”. In such a model, implementing competition between the two mechanisms is a natural way to reach a decision between the two possible outcomes (see also Yuille & Grzywacz, 1998). This would lead to a situation where only one interpretation, integration or segmentation, can prevail at any given moment, as is observed perceptually, but would also offer a natural mechanism for the dynamics of bi-stability we have observed, including the plaid-motion analogue of Levelt's second proposition.

More generally, our findings point to the need to revise how other bi-stable phenomena are understood as well. The perceptual alternations observed for virtually every ambiguous stimulus may reflect not just passive adaptation, but a general strategy adopted by the brain, of implementing mutual inhibition between competing interpretations of sensory stimuli. This idea, which may appear radical at first glance, is in fact consistent with known brain architecture which is rich in reciprocal connections and inhibitory synapses. Research in other domains of perceptual bi-stability may offer new insights on whether this architecture is related to general computational principles.

Notes

1- Under certain conditions, “mixture states” are reported, i.e., percepts which are a superposition of the competing stimuli. But those are fairly rare, e.g., the “fusion” observed at very low contrast in binocular rivalry (Burke, Alais & Wenderoth, 1999, Liu, Tyler & Schor, 1992). The ‘patchy’ percepts sometimes observed in binocular rivalry are not mixture states in this sense, since in every spatial location only one monocular image is perceived.

2- According to this view “mixture states”, which are exceptions to the mutual exclusivity principle, may reflect imperfections in how mutual exclusivity is implemented by the brain.

3- Von Grunau (1993) reported a shortening of the perceptual epochs over time, but this conclusion was most likely due to a methodological problem in how the average durations were computed; see Hupé and Rubin (2003).

4- The first epoch was always coherency only for rectangular-wave plaids. For sinusoidal-gratings plaids, the first percept could be either coherent or transparent and its average duration did not differ from that of subsequent epochs of the same type (see also below, section 8.3, and Hupé and Rubin 2003).

5- Ceiling/floor effects, which are the only sources of deviations from linearity in the effect of alpha, can be observed also with the dynamics approach, they occur, if at all, only at the extremes. For very small Alpha values, coherency is so strong that the entire observation time, albeit long, may still not yield a single transparency epoch. Consequently, $C/[C+T]$ hits the ceiling value of 1, breaking the linear (middle) portion of the curve. Similarly, for very large alpha values transparency may dominate to the extent that $C/[C+T]$ hits the floor value of 0.

6- Subsequent work has shown that Levelt's second proposition is somewhat overstated: in many cases, changing the strength of one eye's stimulus will affect also the dominance durations of that eye (Bossink, Stalmeier & De Weert, 1993). However, it is generally the case that changing stimulus strength in one eye has a much more pronounced effect on the dominance durations of the other eye (Leopold & Logothetis, 1996, Logothetis, Leopold & Sheinberg, 1996, Mueller & Blake, 1989).

References

- Adelson, E.H., & Movshon, J.A. (1982). Phenomenal coherence of moving visual patterns. *Nature*, 300, 523-525.
- Blake, R. (1989). A neural theory of binocular rivalry. *Psychological Review*, 96 (1), 145-167.
- Bossink, C.J., Stalmeier, P.F., & De Weert, C.M. (1993). A test of Levelt's second proposition for binocular rivalry. *Vision Research*, 33 (10), 1413-1419.
- Braddick, O. (1993). Segmentation versus integration in visual motion processing. *Trends in Neurosciences*, 16 (7), 263-268.
- Burke, D., Alais, D., & Wenderoth, P. (1999). Determinants of fusion of dichoptically presented orthogonal gratings. *Perception*, 28 (1), 73-88.
- Hupé, J.M., & Rubin, N. (2003). The dynamics of bi-stable alternation in ambiguous motion displays: a fresh look at plaids. *Vision Res*, 43 (5), 531-548.
- Kim, J., & Wilson, H.R. (1993). Dependence of plaid motion coherence on component grating directions. *Vision Research*, 33 (17), 2479-2489.
- Laing, C.R., & Chow, C.C. (2002). A spiking neuron model for binocular rivalry. *J Comput Neurosci*, 12 (1), 39-53.
- Lehky, S.R. (1988). An astable multivibrator model of binocular rivalry. *Perception*, 17 (2), 215-228.
- Leopold, D.A., & Logothetis, N.K. (1996). Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature*, 379, 549-553.
- Levelt, W.J.M. (1968). On Binocular Rivalry. (p. 107p). The Hague - Paris: Mouton.

- Liden, L., & Pack, C. (1999). The role of terminators and occlusion cues in motion integration and segmentation: a neural network model. *Vision Res*, 39 (19), 3301-3320.
- Liu, L., Tyler, C.W., & Schor, C.M. (1992). Failure of rivalry at low contrast: evidence of a suprathreshold binocular summation process. *Vision Research*, 32 (8), 1471-1479.
- Logothetis, N.K., Leopold, D.A., & Sheinberg, D.L. (1996). What is rivalling during binocular rivalry? *Nature*, 380 (6575), 621-624.
- Movshon, J.A., Adelson, E.H., Gizzi, M.S., & Newsome, W.T. (1985). The analysis of moving visual patterns. In: C. Chagas, R. Gattas, & C. Gross (Eds.), *Pattern Recognition Mechanisms* (pp. 117-151). Rome: Vatican Press.
- Mueller, T.J., & Blake, R. (1989). A fresh look at the temporal dynamics of binocular rivalry. *Biological Cybernetics*, 61 (3), 223-232.
- Von Grunau, M., & Dubé, S. (1993). Ambiguous plaids: switching between coherence and transparency. *Spatial Vision*, 7 (3), 199-211.
- Wallach, H. (1935). *Über visuell wahrgenommene Bewegungsrichtung*. *Psychologische Forschung*, 20, 325-380.
- Wallach, H. (1976). *On Perception*. (p. 490). New York: Quadrangle.
- Wilson, H.R., & Kim, J. (1994). A model for motion coherence and transparency. *Visual Neuroscience*, 11, 1205-1220.
- Wuerger, S., Shapley, R., & Rubin, N. (1996). "On the visually perceived direction of motion" by Hans Wallach: 60 years later. *Perception*, 25 (11), 1317-1367.
- Yuille, A.L., & Grzywacz, N.M. (1998). A theoretical framework for visual motion. In: T. Watanabe (Ed.) *High-level Motion Processing* (pp. 187-211). Cambridge: MIT Press.