

Hierarchical Stochastic Modelling in Multistable Perception

Dissertation

zur Erlangung des akademischen Grades

doctor rerum naturalium

(Dr. rer. nat.)

*genehmigt durch die Fakultät für Naturwissenschaften
der Otto-von-Guericke-Universität Magdeburg*

von Robin Cao, MSc

geboren am 11. November 1987 in Lyon (Fr)

Gutachter: Prof. Jochen Braun, PhD

Assoc. Prof. Rafal Bogacz, PhD

eingereicht am: 30. November 2016

verteidigt am: 20. Juni 2017

Zusammenfassung

Multistabile Wahrnehmung entsteht bei der Betrachtung mehrdeutiger Bilder, wenn gegenstzliche Wahrnehmungen des Bildes spontan alternieren. Trotz gleichbleibender Stimulation schlt die Wahrnehmung unweigerlich frher oder spter um. Es wird vermutet, dass diese seltsame Entkopplung zwischen Stimulation und bewusster Wahrnehmung tiefe Einblicke in die neuronalen Ablufe geben knnte, mit denen das Gehirn sensorische Information ver- und entschlsselt. Trotz erheblicher experimenteller und theoretischer Fortschritte sind die Mechanismen, die der multistabilen Wahrnehmung zugrundliegen, noch nicht vollstndig verstanden.

Bisherige Erklrungsanstze haben sich auf den Umschaltvorgang konzentriert und mehrere plausible dynamische Mechanismen vorgeschlagen, unterliegen jedoch wichtigen Beschrnkungen. Insbesondere erfllen sie die Skalierungseigenschaft der multistabilen Wahrnehmung nicht, ein allgemeingltiger und erstaunlich robuster experimenteller Befund ber die statistische Verteilung der Umschaltzeiten. Obwohl die durchschnittlichen Umschaltzeiten verschiedener Betrachter und verschiedener Bilder sich stark unterscheiden, bleibt die Form der Verteilungsdichte unverndert. Die vorliegende Dissertation schlt einen neuartigen Ansatz zur Erklrung multistabiler Wahrnehmung vor, welcher diese einzigartige Skalierungseigenschaft in das Zentrum der Betrachtung rckt.

Ganz allgemein verraten die zufllig verteilten Zeitpunkte von Wahrnehmungsentscheidungen viel ber die zugrundelegende Dynamik neuronaler Aktivitt. Im Zusammenhang mit Detektions- und Diskriminationsleis-

tungen, kann der Zeitpunkt der Wahrnehmungsentscheidung als langsame und rauschbehaftete Ansammlung sensorischer Informationen bis zum Erreichen einer Wahrnehmungsschwelle beschrieben werden. Die Einföhrung dieses first-passage-time Ansatzes hat zahlreiche wichtige Aspekte der neuronalen Ablufe enthllt, denen Wahrnehmungsentscheidungen zugrundeliegen. Die vorliegende Dissertation bertrgt diesen Ansatz auf den Fall der multistabilen Wahrnehmung und die Untersuchung ihrer charakteristischen Skalierungseigenschaft, indem sie die Zeitpunkte des Umschaltens als first-passage-time Problem behandelt. Auf dem Wege einer genauen, vergleichenden Analyse unterschiedlicher Zufallsablufe werden die Ursachen der Skalierungseigenschaft der multistabilen Wahrnehmung erkannt und beschrieben.

Die Auswirkungen dieser Einsichten werden dann benutzt, um einen neuartigen, hierarchischen Mechanismus der multistabilen Wahrnehmung vorzuschlagen, dessen Wechselwirkungen sowohl top-down als auch bottom-up verlaufen. Whrend anfngliche Wahrnehmungsentscheidungen hauptsächlich von sensorischen Reizen, durch bottom-up Wirkungen, bestimmt sind, wird der Umschaltvorgang von dem vorherrschenden Wahrnehmungszustand, durch top-down Wirkungen, ausgelst. Das sich daraus ergebende Modell erklrt die verfügbaren experimentellen Beobachtungen ber multistabile Wahrnehmung sowohl quantitativ als auch qualitativ umfassend und mit einem bislang unerreichten Grad der Genauigkeit. Des weiteren schliet die hierarchische Wechselwirkung, welche in der vorliegenden Dissertation beschrieben ist, auch einige wichtige Aspekte anderer Theorien und Modelle mit ein, welche die Wahrnehmungsentscheidungen und kortikale Informationsverarbeitung ganzheitlich betrachten. Aus diesem Grund knnte die beschriebene hierarchische Wechselwirkungen ber die multistabile Wahrnehmung hinaus auch fr weitere kognitive Phnomene von Bedeutung sein.

Abstract

Multistable perception is observed during the perception of ambiguous figures, where spontaneous alternations between alternative perceptions of the figures occur. Although the stimulus does not change, perception eventually and ineluctably reverses. It is thought that understanding this peculiar dissociation between incoming sensory input and its conscious representation, may provide considerable insights into the neural processes by which the brain encodes and decodes information. Despite extensive experimental and theoretical advances, the mechanisms underlying perceptual reversals are not fully understood.

Traditional modelling approaches have focused on characterizing this reversal behaviour and have proposed several plausible underlying mechanisms, but suffer from important limitations. In particular, they do not satisfy the scaling property of multistable perception, a general and strikingly robust empirical aspect of reversal timing. Although average reversal times widely vary between observers and experimental conditions, the shape of reversal times densities remains invariant. This thesis proposes a novel approach to multistable perception that focuses on this unique scaling property.

More generally, the stochastic timing of perceptual choice events reveals much about the underlying dynamics of neural activity. In the context of detection and discrimination tasks, choice events can be described as a gradual and noisy accumulation of sensory information to a perceptual threshold. The introduction of this ‘first-passage-time’ framework has uncovered important aspects of neural mecha-

nisms implementing perceptual decisions. This study adapts this approach to the case of multistable perception, and to the study of its characteristic scaling property, consider the timing of perceptual reversals from the perspective of a first-passage-time problem. By means of a detailed comparative analysis of several random walk models, the origins of the scaling property of multistable perception are identified and discussed.

The implications of these results are then exploited to introduce a new hierarchical mechanism of perceptual reversals, which features both bottom-up and top-down interactions. While initial perceptual choices are mainly driven by the sensory input, via bottom-up influences, reversal dynamics is driven by the current perceptual state, via top-down influences. The resulting model accounts for the available behavioural data, qualitatively and quantitatively to a level of detail which had not previously been achieved. Additionally, the hierarchical dynamics described in this thesis rediscovers some important aspects of decision-making theories and of other large-scale cortical models. Thus, the proposed hierarchical dynamics may carry wider implications well beyond the study of multistable phenomena.

Erklärung

Hiermit erkläre ich, dass ich die von mir eingereichte Dissertation zum dem Thema:

Hierarchical Stochastic Modelling in Multistable Perception

selbständig verfasst, nicht schon als Dissertation verwendet habe und die benutzten Hilfsmittel und Quellen vollständig angegeben wurden.

Weiterhin erkläre ich, dass ich weder diese noch eine andere Arbeit zur Erlangung des akademischen Grades doctor rerum naturalium (Dr. rer. nat.) an anderen Einrichtungen eingereicht habe.

oder

entsprechend abändern

Bristol, Grossbritannien, den 29. November 2016

Robin Cao, MSc

Acknowledgements

First of all, I would like to thank my supervisors, Prof. Jochen Braun and Maurizio Mattia, for guiding me through this scientific journey, and offering me a mentorship from which I have learnt so much more than academic knowledge.

To my mom, my nan and my brother, for existing at all.

To my colleagues in Magdeburg and Rome, and in particular, to Patrick Camilleri, whose departure from KoBi left a great void, and to Christoph Bauermeister for the many cigarette-break discussions.

To all my friends and ‘crews’ across Europe, in Lyon, Paris, Bristol, Magdeburg and Rome, you should all know who you are by now.

A gert massive thank you to Natalie for her resilience in enduring the worst of my moods and for believing in me when I did not.

This work is dedicated to my dad, who I like to think would have been rather proud.

Contents

1	Introduction	1
1.1	Perception and uncertainty	1
1.2	Sources of variability	2
1.3	Why study multistable perception?	4
1.4	Current models of multistable perception	7
1.5	Thesis outline	9
2	Scaling property of multistable perception	15
2.1	Behavioural approach	16
2.1.1	Invariance of shape	16
2.1.2	Invariance of moments	17
2.1.3	Experimental methods	20
2.1.4	Behavioural observations	22
2.2	Modelling approach	25
2.2.1	First-passage-time framework	25
2.2.2	Comparison criterion	28
2.2.3	Sufficient condition for the scaling property	32
2.3	Gaussian diffusion-to-bound	34
2.3.1	Standard Wiener process	36
2.3.2	Modified Wiener Process	39
2.3.3	Standard Ornstein-Uhlenbeck Process	40
2.3.4	Modified Ornstein-Uhlenbeck Process	42
2.4	Summary	43

3	The Ehrenfest Urn Model	47
3.1	Introduction	48
3.2	Ehrenfest urn model	49
3.2.1	Definition	49
3.2.2	Detailed balance	51
3.2.3	Mean and variance of the Ehrenfest process	52
3.2.4	General Solution	54
3.3	Moments of the first-passage-time density	57
3.3.1	First-passage-time density	57
3.3.2	Moments formulae	58
3.3.3	Simulations	61
3.3.4	Low-Threshold Approximation	65
3.4	Regime Identification	65
3.4.1	Drift- and noise-dominated regimens and the scaling property	65
3.4.2	Drift- and noise-dominated regimens and the distribution shape	66
3.4.3	Interacting pool	69
3.5	Summary	73
4	Scaling Property and Distribution Shape	75
4.1	Generalized Ehrenfest process	76
4.2	Balanced Poisson process	78
4.2.1	Definition and Master Equation	78
4.2.2	Exact solution	79
4.2.3	First-passage-time distribution	80
4.3	Cox-Ingersoll-Ross process	82
4.3.1	System-size expansion	82
4.3.2	Scaling property in the CIR process	84
4.4	Distinguishing properties of the GE process	87
4.4.1	Overview	87
4.4.2	Granularity, finite-size effects	89
4.4.3	Asymmetric dispersion far-from-equilibrium	92
4.5	Neurally plausible realization of a GE process	96
4.6	Summary	100

5	Hierarchical Multistable Dynamics	103
5.1	Introduction	104
5.2	Competition, adaptation and noise	107
5.3	Qualitative hierarchical dynamics	109
5.3.1	Levels of hierarchy	111
5.3.2	Interactions within hierarchy	113
5.3.3	Qualitative description of dynamical behaviour	115
5.4	Formal hierarchical dynamics	121
5.4.1	Dynamical equations	121
5.4.2	Analysis of decision level: perceptual threshold and switching behaviour	124
5.4.3	Analysis of evidence level: evidence-bias	126
5.5	Agreement with behavioural data	129
5.5.1	Model optimization	129
5.5.2	Fit results: average dominance durations	132
5.5.3	Levelt's Proposition I and II	132
5.5.4	Levelt's Proposition IV	136
5.5.5	Fit results: scaling property and distribution shape	140
5.6	Non-stationary dynamics	143
5.6.1	Correlations and structure of correlations	144
5.6.2	Slow fluctuating drift of sensory evidences	149
5.7	Summary	151
6	Conclusions	153
	Appendices	167
A	Appendix A	169
B	Appendix B	177
C	Appendix C	181

List of Figures

2.1	Shape invariance of reversal time densities	17
2.2	Examples of canonical multistable displays	20
2.3	Scaling property of multistable perception	23
2.4	First-passage-time scenario	27
2.5	Drift- and noise-dominated regimes	30
2.6	Diffusive processes	35
2.7	Moments of the Wiener process' FPT density	37
2.8	Scaling property in the Wiener process	38
2.9	Moments of the modified Wiener process' FPT density	39
2.10	Scaling property in the modified Wiener process	40
2.11	Moments of the OU process' FPT density	41
2.12	Scaling property in the OU process	42
2.13	Moments of the modified OU process' FPT density	43
2.14	Scaling property in the modified OU process	44
2.15	Schematic mechanism for the scaling property	45
2.16	Schematic mechanism for the scaling property	46
3.1	Local bistable attractors	50
3.2	Theory vs. Simulations (CV)	62
3.3	Theory vs. Simulations (Skewness)	63
3.4	Theory vs. Simulations (Kurtosis)	64
3.5	Regime identification (CV)	67
3.6	Regime identification (Distribution shape)	68
3.7	Regime identification (CV and Skewness)	68
3.8	Scaling property in an assembly of interacting bistable units	71

3.9	Dynamics in an assembly of coupled bistable units	72
4.1	Moments of the Ehrenfest process' FPT density	77
4.2	Scaling property in the Ehrenfest process	77
4.3	Moments of the balanced Poisson process' FPT density	81
4.4	Scaling property in the balanced Poisson process	82
4.5	Moments of the CIR process' FPT density	86
4.6	Scaling property in the CIR process	86
4.7	Increment and decrement rates in the Ehrenfest process	90
4.8	Input- and state-dependent drift and variance in the Ehrenfest process	91
4.9	Asymmetric diffusion of the Ehrenfest process	94
4.10	Symmetric diffusion of the balanced Poisson process	95
4.11	Spiking neurons simulations	98
4.12	Spiking neurons simulations	101
5.1	Moments of the escape process' FPT density	108
5.2	Scaling property in the escape model	109
5.3	Three-levels hierarchy of perceptual representations	112
5.4	Mechanism for perceptual selection	115
5.5	Mechanism for perceptual reversal	117
5.6	Joint habituation/recovery dynamics of sensory evidences	120
5.7	Structure of the hierarchical model	121
5.8	Perceptual threshold	125
5.9	Evidence layer dynamics and contrast modulations	130
5.10	Model fit to average dominance durations	133
5.11	Mechanism of Levelt's 2nd proposition	135
5.12	Mechanisms of Levelt's 4th proposition: threshold modulations . . .	138
5.13	Mechanisms of Levelt's 4th proposition: amplitude of threshold mod- ulations	139
5.14	Model fit to CV and Skewness	141
5.15	Correlation of alternation sequences	145
5.16	Structure of correlations	147
5.17	Sources of correlations	149
5.18	Sensory origins of correlations	151
A.1	Equilibrium distribution of the GE and CIR processes	174

B.1	Optimization results	178
B.2	Optimal parameters inspection	179
C.1	Transition-triggered average (low contrasts)	182
C.2	Transition-triggered average (high contrasts)	183
C.3	Transition-triggered average (unbalanced, low-high)	184

1.1 Perception and uncertainty

Our perception of the world relies on the translation of its physical attributes into coherent and reliable mental representations. This implies, on the one hand, the ability to form internal representations associated with the sensory input (encoding), and on the other hand the ability to access these internal states by extracting relevant information from the sensory input (decoding), which we progressively learn from birth by interacting with our environment. In the context of perceptual decision-making, the joint operation of both encoding and decoding processes is necessary: our perception is often challenged by various forms of uncertainty, for instance, a lack of prior knowledge about incoming information, a lack of incoming information itself, or both. Sensory input is intrinsically ambiguous; prior knowledge about the properties of the world is always needed to allow incoming stimuli to be disambiguated and appropriately categorised. One could in principle interpret the back and forth motion of an object as a deformation of the object itself, as if it were gradually shrinking and growing, and conversely. The prior knowledge, acquired through perceptual learning, that some objects are not likely to shrink or grow, would help in resolving this specific ambiguity. The most plausible alternative would then be uniquely selected, and other alternatives, dismissed [von Helmholtz, 1866, Barlow, 1990, Gregory, 2009].

Object recognition, for instance, requires knowledge about its physical attributes,

which may themselves widely differ between objects with the same function. It also requires knowledge about alternative objects, which may share common characteristic features, but differ in functions. Correct recognition of an object, and more generally, correct interpretation of natural scenes, relies on the ability of our visual system to cope with these different sources of uncertainty, in order to propose a stable and reliable perceptual representation in awareness at all times. However, in situations where incoming stimuli are corrupted by noise, inconsistencies or ever-changing qualities, and may evoke alternative percepts, even extensive prior knowledge about all possible interpretations may not be sufficient to make a reliable perceptual decision.

1.2 Sources of variability

In addition to incoming sensory information being ambiguous and fluctuating, neural responses at all levels of the visual processing chain, from the sensory level, to higher stages of cognition, are also inherently noisy [Arieli et al., 1996, Tsodyks et al., 1999]. Perception itself is stochastic: sensory inputs are not always faithfully interpreted and perceptual decisions may have probabilistic outcomes. This is well captured by the application of signal-detection theory to the psychophysics of detection and discrimination tasks, in particular, when stimuli are weak and close to detection thresholds [Swets, 1964, Green and Swets, 1966]. At the level of single-neurons, the variability in the spiking activity of a single-neuron depends on fluctuations of its ion-channels, which condition the release of action potentials (also termed ‘channel noise’ [White et al., 2000]), and on the synaptic input it receives from other neurons (also termed ‘synaptic noise’ [Destexhe and Paré, 1999]). Individual variability is then carried upwards to the level of neuron populations, depending on the structure of interactions and correlations between the neurons populating a given assembly, as well as on its size [Brunel and Hakim, 1999, Mattia and Del Giudice, 2002, Faisal et al., 2008]. Also, when considering neural assemblies on increasingly large scales, fluctuations do not only build-up in a bottom-up manner as a result of neural variability, but may also be affected

by top-down influences, such as fluctuations in attention, or in other cognitive states.

The relative contributions to behavioural variability of noise from external or sensory sources, are not well established. This constitutes a highly active area of research, to which the present study provides additional contributions [Gold and Shadlen, 2007, Ratcliff and McKoon, 2008, Nienborg and Cumming, 2009, Amitay et al., 2013, Litwin-Kumar and Doiron, 2012, Wimmer et al., 2015]. While the nature of processes involved in brain dynamics ineluctably results in noisy neural and behavioural responses, such variability could in fact be an important functional feature of neural computations underlying decision-making [Bogacz et al., 2006, Ma et al., 2006, Beck et al., 2008, Fiser et al., 2010, Churchland et al., 2011, Pouget et al., 2013, Haefner et al., 2016]. Stochastic dynamics may be desirable for perceptual decisions to select the representation which best accounts for the available sensory evidence, while simultaneously allowing for other alternatives to be considered and even reversed to. If additional evidence is made available which make the current interpretation of the stimulus somewhat improper, maintaining this interpretation rather than switching to a more likely alternative - by ignoring new sensory information or exclusively focusing on a subset of sensory information - could have dramatic consequences.

Rather than simply interfering with neural computations, stochastic variations of activity may act to appropriately balance the relative contributions of prior knowledge and current sensory input to decision-making. Performances in the categorisation of sensory inputs would sensibly differ if one were to arbitrarily outweigh the other, irrespectively of the nature of the stimulus. Perception is most stable and least sensitive to changes in the environment when overly relying on prior knowledge: even compelling environmental changes may not be reflected at the perceptual level [Kersten et al., 2004, Pastukhov et al., 2013]. For instance, exotic stimuli may be grossly misinterpreted on the basis of a slight similarity to a familiar stimulus, rather than in terms of a new category. Conversely, perception is most sensitive and least stable when overly relying on current sensory information: even unconvincing environmental changes may be unduly reflected at the perceptual level. This could

cause known stimulus to be interpreted in terms of a new category, rather than in a known category, on the basis of slight dissimilarities with familiar stimuli.

This trade-off between perceptual stability and sensitivity is reminiscent of the ‘exploration-exploitation’ dilemma, originally formulated in the context of reinforcement learning [Sutton and Barto, 1998]. For humans, animals or even commercial organisations, survival and prosper development requires them to adapt to the many unknowns and ever-changing properties of their environment. In a foraging scenario, an exploitation-only strategy may result in the depletion of a given resource, while an exploration-only strategy may result in insufficient accumulation of resources, neither of which would be beneficial to thrive in the long term. Equivalently, optimal perceptual decisions in challenging conditions would require both strategies to be adequately balanced, to ensure that changes in our perception can steadily match changes in our environment.

1.3 Why study multistable perception?

Despite the seemingly artificial aspects of its practical implementation, which involve the use of specific and somewhat unnatural stimuli, multistable phenomena are thought to open a particularly convenient window on the stochastic and exploratory dynamics of perceptual inference [Leopold and Logothetis, 1999, Blake and Logothetis, 2002, Sterzer et al., 2009, Pastukhov et al., 2013]. This is precisely because even when stimulus qualities of multistable displays do not change, perceptual experience does in a particularly prominent fashion. This way, the mechanisms underlying visual inference may be approached in a particular case, without having to consider the more general and considerably more complex setting described previously, where perception is modulated by an ever-changing environment.

Multistable stimuli are such that alternative interpretations of the stimulus interchange in time, as they can not be disambiguated. Perceptual choices are thus continually reconsidered, remain ultimately inconclusive, and as a result, subjects experience reportable switches in awareness. Importantly, reversals in perceptual dominance are spontaneous: although volitional control can affect the duration

of stable appearances, reversals ineluctably occur [von Helmholtz, 1866, Pastukhov and Braun, 2007, Kornmeier et al., 2009]. Multistable situations may reflect the balance between ‘exploration’ and ‘exploitation’, spontaneously positioning perceptual dynamics in a self-organised regime, optimally balancing stability and sensitivity. This allows fluctuations to probe whether all possible interpretations of the stimulus have been considered (exploration), while maintaining the current interpretation stable (exploitation) [Kim et al., 2006, Pastukhov et al., 2013].

This suggests that unresolved ambiguities prompts sensory evidences to be continuously re-evaluated, and confronted to the current state of awareness. Since all sensory inputs are intrinsically ambiguous, multistability may be a hallmark of perceptual inference, and also occur during the interpretation of natural scenes. One characteristic of multistable displays is that they evoke substantially distinct states in prior knowledge, making transitions between those states particularly sharp and noticeable. Presumably, natural displays may evoke a more continuous and high-dimensional set of alternative states, so that perceptual transitions would be comparatively smooth and unnoticeable.

In addition to such puzzling reversal behaviour, the temporal dynamics of multistable displays is characterised by a vast array of experimental observations which reveal much about the timing of underlying neural processes; they have been and still are intensively discussed, but are not fully accounted for. Stable dominance periods between reversals (or dominance durations) obey striking statistical properties, highly consistent across observers, display types and even for other sensory modalities (e.g. for audition [Pressnitzer and Hupe, 2006, Winkler et al., 2012] and touch [Carter et al., 2008]). The apparent generality of these findings further supports the idea that the study of multistability can unveil fundamental insights into the mechanisms of perceptual inference.

Our study will focus on the following experimental observations:

- **Reversal behaviour:** competing percepts are *mutually-exclusive*, reversals are *spontaneous* and *stochastic* [Leopold and Logothetis, 1999].
- **Distribution of dominance durations:** for fixed stimulus qualities, dominance durations are *Gamma-distributed*, with coefficient of variation $c_v \approx 0.6$, and skewness $\gamma_1 \approx 2c_v$ [Cao et al., 2016].
- **Levelt's propositions:** average dominance durations can be modulated by variation of input-levels, which can be summarized in two main characteristic laws. First, relative variations in stimuli strength cause non-linear and asymmetric modulations of dominance durations. In particular, stronger percepts are more noticeably affected by such changes than weaker percepts (originally known as Levelt's second proposition). Second, greater absolute stimulus strengths yields greater alternation rate (originally known as Levelt's fourth proposition) [Levelt, 1965, Klink et al., 2008, Brascamp et al., 2015].
- **Scaling property:** the distribution dominance durations consistently conserves its shape for different displays, observers and stimulus qualities, while average dominance durations may considerably vary. This implies that the coefficient of variation and skewness of dominance duration remain essentially constant despite large variations of the mean.
- **History-dependence:** reversal sequences show weak but consistent *sequential correlation* between successive dominance durations [van Ee, 2009]. Longer-lasting correlations are also revealed in new experimental data by the heterogeneity of reversal counts in different time windows (*burstiness* of reversal sequences) [Cao et al., 2015a,b].

1.4 Current models of multistable perception

Current models of multistable perception reproduce several important aspects of multistable perception, by employing neurophysiologically plausible assumptions to generate a noisy attractor-like dynamics. Such models detail the possible mechanisms underlying perceptual reversals, usually by means of reduced ‘mean-field’ descriptions [Lehky, 1988, Wilson, 2007, Shpiro et al., 2007, Moreno-Bote et al., 2010]. More detailed implementations using networks of spiking neurons have also been proposed [Laing and Chow, 2002, Moreno-Bote et al., 2007]. In general, current models rely on the following: local recurrent excitation, global mutual-inhibition, exogenous fluctuations, and neural adaptation. Common adaptation mechanisms in the models are usually related to synaptic dynamics (e.g. synaptic depression), or channel dynamics (e.g. spike-frequency adaptation). Perceptual outcomes are determined by the global activity of self-sustaining stable states (attractors) obtained via recurrent excitation, where mutual-inhibition between competing attractor states ensures the exclusivity of the dominant perceptual appearance [Amit and Brunel, 1997]. Spontaneous transitions between attractor states, implementing perceptual reversals can be obtained by adding a sufficient amount of noise, or a sufficient amount of adaptation [Laing and Chow, 2002, Moreno-Bote et al., 2007, Shpiro et al., 2009].

Both alternative mechanisms can be used independently to account for reversal behaviour. However they must be appropriately balanced to reproduce the observed distribution of dominance durations, but also the sequential correlation between successive dominance periods. In particular, strong adaptation typically produces reversal sequences with significantly greater correlation than those observed experimentally [Lehky, 1988, Shpiro et al., 2009]. A regime where perceptual reversals are mainly noise-driven as been reported to provide both realistic distributions dominance and weak correlations in alternation sequences [Brascamp et al., 2006, Moreno-Bote et al., 2007, Shpiro et al., 2009]. This regime is also consistent the results from perturbation experiments, which suggest a gradual weakening of the dominant visual appearance between reversals [Wolfe,

1984, Nawrot and Blake, 1989, Petersik, 2002, Kang and Blake, 2010].

Note that although adaptation must be weak, it still plays an important role. If noise were the sole driving force behind perceptual alternations, reversals statistics would reduce to the well-known Kramer escape problem and dominance durations would be exponentially distributed [Kramers, 1940, Moreno-Bote et al., 2007, Shpiro et al., 2009, Gigante et al., 2009]. In this regime, however, reversals are unlikely to occur before the system has started to adapt, for instance, in the moments directly following the precedent reversal. Adaptation gradually destabilize dominant and suppressed states, so that after some time, noise is able to initiate the next perceptual switch. This can be seen as a refractory period, which effectively limits the amount of extremely short reversal times, preventing dominance durations to distribute exponentially.

The characteristic modulations of average dominance durations under variations of stimulus strength (Levelt's propositions), are perhaps the most commonly used criterion to ascertain the validity of computational models of multistability. Since Levelt's original work, these propositions have been extensively studied, both experimentally and theoretically, and have been subjected to several reformulations as new evidence was made available [Levelt, 1965, Klink et al., 2008, Kang, 2009, Brascamp et al., 2015]. In some instances, previously proposed models present a good agreement with Levelt's second proposition (see for instance [Laing and Chow, 2002, Wilson, 2007, Moreno-Bote et al., 2007]). In their studies, Laing and Wilson established approximate formulas for the mean duration of dominance periods, which explicitly depend on input-levels. These formulas provide an important basis for the qualitative understanding of Levelt's second proposition, in terms of the characteristic profile of the adaptation mechanism shared by these models [Brascamp et al., 2015]. However, it is important to note that the agreement of these models with Levelt's propositions is, in general, partial. They typically explored restricted range of stimulus strength combinations, with only one of two applied stimulus strengths being modulated.

Also, they have encountered difficulties in reproducing Levelt's fourth proposi-

tion. Contrary to experimental evidences, the reversal rate in adaptation-based models is not a monotonic function of absolute stimulus strength, a result which has prompted several detailed studies, without being confirmed experimentally [Shapiro et al., 2007, Curtu et al., 2008, Seely and Chow, 2011, Brascamp et al., 2015].

An important limitation adaptation-based models is that they can not conciliate modulations of average dominance durations with higher-order statistical properties of reversal times sequences (i.e. the scaling property). Observed reversal sequences obey consistent and striking statistical properties: mean dominance durations can span across a few orders of magnitude (from the hundreds millisecond to nearly 100 seconds) depending on subjects, displays types and stimulus levels [Fox and Herrmann, 1967, Borsellino et al., 1972, Walker, 1975, Zhou et al., 2004, Brascamp et al., 2005, Cao et al., 2016], higher-order moments of dominance durations distributions scale proportionally with the mean. The distribution shape is thus invariant and remains close to a gamma distribution (i.e. the coefficient of variation $c_v \in [0.4, 0.6]$ and the skewness $\gamma_1 \approx 2c_v$) [Levelt, 1967, Blake et al., 1971, Walker, 1975, De Marco et al., 1977, Murata et al., 2003, Pastukhov and Braun, 2007, Cao et al., 2016]. Also, sequential correlations between successive dominance durations are generally non-significant, and at best, small (around 0.1, 0.2) [Fox and Herrmann, 1967, Lehky, 1995, van Ee, 2009, Pastukhov et al., 2013].

As previously mentioned, in order to obtain the coefficient of variation obtained experimentally, along with the weak correlation, adaptation-based models must precisely balance the contributions of noise and adaptation to reversal dynamics [Shapiro et al., 2009]. However, they can only maintain this balance over a narrow regime, and therefore can only reproduce Levelt's propositions at the expense of the scaling property. Ideally, both results should be simultaneously verified.

1.5 Thesis outline

In this thesis, I present a new model of multistable perception, which resolves several of the issues encountered by other models, by approaching the relationship

between neural and behavioural fluctuations in a more comprehensive manner. The results presented here are articulated around two peer-reviewed publications, and a manuscript in preparation [Cao et al., 2014, 2016]. The originality of our approach is that we considered the possibility that higher-order statistical properties, such as the scaling property of multistable perception, may provide important constraints on the dynamics of the underlying neural systems. Furthermore, that these properties may in turn provide additional constraints to design a perceptual reversals mechanism, on the grounds that the adequate mechanism should preserve such properties. Traditional models of multistable perception usually consider a deterministic mechanism of perceptual reversals, to which external noise is added and tuned to match the variability of behavioural responses. Instead, we constrained the deterministic properties of our model by first examining the properties of its fluctuations.

A common approach in the study of perceptual decisions is to represent the timing of neural events as the time-to-threshold of a stochastic decision variable. It is assumed that mechanisms underlying perceptual decisions can be reduced to a random walk model, representing a gradual and noisy accumulation of sensory evidences, evoked by the presentation of a stimulus. Perceptual choices occur when the decision variable reaches a threshold level. These so-called *diffusion-to-bound* models have been widely used in the modelling of two-alternatives forced choice tasks (2AFC), a well-studied paradigm of perceptual decision-making [Ratcliff and Smith, 2004, Smith and Ratcliff, 2009]. They explain response time distributions for both correct and error responses as well as their dependence on motion strength, allowing for a direct link between neural and behavioural dynamics to be drawn. This approach is supported by a large body of neurophysiological studies which have identified cortical areas predictive of perceptual choices and showing a qualitatively similar dynamics as predicted theoretically (for instance, in parietal and prefrontal cortical area [Kim and Shadlen, 1999, Schall, 2001, Sugrue and Corrado, 2005, Gold and Shadlen, 2007, Shadlen and Kiani, 2013]).

The present study sought to extend this stochastic accumulation framework to the case of multistable phenomena. The peculiar dissociation between stimulus and

perceptual experience observed in multistable perception, along with the seemingly random nature of reversal sequences, suggests the existence of a stochastic decision variable driving perceptual dynamics from behind the scenes. Also, cortical sites associated with, respectively, accumulating and evaluating competing evidence for different forms of multistable displays appear to overlap with the sites identified for 2AFC tasks [Roitman and Shadlen, 2002, Knapen et al., 2011, Shadlen and Kiani, 2013]. This supports the idea that 2AFC tasks and multistable perception may to some extent share common neural underpinnings, and further justifies an approach in terms of a diffusion-to-bound model, to characterise the timing of neural events underlying perceptual reversals.

From a modelling perspective, statistical physics and the theory of stochastic processes constitute a natural and powerful way to bridge the gap between the fluctuating dynamics of neural systems and the fluctuating dynamics of behavioural responses. As a first step, we performed a detailed analysis of *first-passage-time* properties for several random walk models, to determine which of these models could satisfy the scaling property, and importantly, which dynamical features allowed them to.

These results are presented in **Chapter 2**. We present our behavioural observations, as well as the minimum conditions for this property to be satisfied by diffusion-to-bound models. The first-passage-time properties of several traditional random walk models are analysed, and their ability to satisfy the scaling property, discussed. We found that to obtain first-passage-time densities which widely differ in their mean but conserve their shape as the stimulus is varied, the noise (or diffusion coefficient) must obey a specific dependence on input-levels [Cao et al., 2016].

Then, we envisaged stochastic accumulation of activity as originating from the collective activation in a finite ensemble of *independent* bistable units. Each unit can become spontaneously active (or inactive), transition rates are modulated by input-levels and occur according to a Poisson process. These discrete neural units are thought to idealize the dynamics of discrete neural nodes such as cortical

columns [Amit, 1995, Mattia et al., 2013]. Collective accumulation of activity in an assembly of such bistable units realizes a birth-death process known as the generalized Ehrenfest urn model. In this case, both mean and variability of the collective dynamics naturally depend on input-levels, and share a common physical origin, so that the scaling property is automatically satisfied.

These results are presented in **Chapter 3**, where we formally introduce the generalized Ehrenfest process and derive several important analytical results [Cao et al., 2014]. In particular, a recursive expression for the moments of the first-passage-time density is obtained, the four first moments are explicitly derived and compared to direct simulations. A regime identification is performed and reveals a specific regime in which the Ehrenfest process reproduces the scaling property as well as the characteristic Gamma-like shape of reversal times densities observed in experiments. Importantly, we extend these results to the case of an assembly *interacting* bistable units, corresponding to the dynamics of a multi-modular cortical network, or cluster of interacting cortical columns.

We originally introduced the Ehrenfest urn model because of its natural disposition to satisfy the scaling property. This is because when collectively accumulating activity with bistable units obeying a Poisson statistics, both deterministic and stochastic components of the dynamics have the same physical origin and appropriately scale with input-level. Unexpectedly, we also found that it could reproduce the characteristic distribution shape of observed reversal time densities. In fact, the reasons underlying this result are not as straightforward.

In **Chapter 4** we sought to clarify this issue by singling out the dynamical properties expressed by the Ehrenfest process which may be responsible for the shape of its first-passage-time distribution. To this aim, we draw an additional comparison with the statistics of two closely-related random walks. While sharing some important characteristics of the Ehrenfest process, in that they also both naturally satisfy the scaling property, they do not reproduce the desired distribution shape, because both processes diffuse symmetrically. In contrast, when operating far-from-equilibrium, the Ehrenfest process is non-Gaussian and diffuses

asymmetrically. In addition, we propose a neurally plausible implementation of the Ehrenfest process by a network of spiking-neurons.

Finally, we redesigned the mechanism underlying perceptual reversals, to allow the resulting model to simultaneously reproduce the scaling property and Levelt's propositions. Similarly to previously proposed models, the activity of mutually inhibited attractor states represents perceptual outcomes, and reversals are driven by the joint action of a gradual destabilisation mechanism of the dominant percept, and by noise. The critical difference with prior models is that we replaced the combination of exogenous noise and deterministic self-adaptation by a discrete random walk (the Ehrenfest process).

In a previous Chapter, we establish that this process could reproduce important statistical properties of multistable perception, thanks to several key dynamical features. To preserve these important features, we implemented perceptual in a hierarchical manner, combining slow random walk, with fast winner-take-all attractor dynamics [Cao et al., 2015a,b]. The former models the accumulating dynamics of subconscious 'evidence' representation at the sensory level, and the latter models the dynamics of conscious 'decision' representation at the behavioural level. Reversals occur as a result of the joint dynamics between the two levels, when the relative strength of mounting evidence sufficiently contradicts the current perceptual state. Importantly, this sudden change of decision at the decision level directly impacts the dynamics of the lower sensory level, reverting the relative accumulation of evidence, and causing sensory populations to reorganise accordingly. Evidence supporting the newly dominant state weakens, eventually contradicting the decision and prompting another reversal.

Although akin to self-adaptation in its effects, our proposed mechanism reflects a joint dynamics of sensory accumulation (supporting the suppressed percept) *driving* the behavioural level via bottom-up projections, and of sensory habituation (supporting the dominance percept) *driven by* the behavioural level via top-down projections.

In **Chapter 5**, we present the detailed implementation of our hierarchical model

of multistable perception. We illustrate the dynamics at each level the hierarchy separately, before gradually introducing the specifics of our reversal mechanism, as well as the governing equations of the model. The model is then fitted to behavioural data up to third-order (mean, variance, and skewness). In particular, the model provide simple mechanisms explaining all of Levelt's propositions accurately, along with the scaling property and shape of reversal time densities.

The structure and dynamics of our model is inspired and supported by the idea that multistable appearances may reflect hierarchical interactions between visual sensory areas encoding low-level stimulus informations, and high-level areas (non necessarily visual) encoding high-level stimulus information [von Helmholtz, 1866, Leopold and Logothetis, 1999, Sterzer and Rees, 2008, Hohwy et al., 2008, Gigante et al., 2009, Kang and Blake, 2010].

Beside qualitatively and quantitatively accounting for experimental data to a level of detail which had not previously been achieved, our model will propose several important implications. Firstly, it predicts history-dependence effects on multiple time-scales in agreement with new behavioural observations. Secondly, it relies on neurophysiologically plausible assumptions which are consistent with observations in non-human primates. Lastly, its operating regime reflects a dynamical mechanisms which may be relevant to the study of perceptual inference in a more general setting than multistable perception.

To conclude, **Chapter 6** will relate several aspects of the model to their possible neurophysiological substrates, and examine its similarities with other important models of perceptual decision-making. We will discuss the possible implications this model may provide in the more general context of perceptual inference.

Scaling property of multistable perception

From an experimental point of view, the scaling property of reversal time densities is known as a general and well established feature of all multistable phenomena, in that it consistently pertains across observers, display types and stimulus qualities. From a modelling point of view, however, this astonishing result has generally been neglected, and remains largely unaccounted for. In an attempt to bridge this gap in knowledge, we followed the intuition that the scaling property may provide important clues on the qualitative nature and dynamical properties of neural structures underlying multistable perception.

Here, we present a simplified approach to this issue in terms of a first-passage-time problem, by assuming that the time between perceptual events reflects the stochastic accumulation of activity in a neural population. First-passage-time statistics can be modulated by changes in the stimulus strength and the scaling property may or may not be satisfied. In the case of traditional drift-diffusion models, a sufficient condition for the scaling property to hold is simply that dispersion rate (i.e. the rate at which activity disperses) and drift rate (i.e. the rate at which average activity accumulates) must remain proportional as stimulus strength vary.

2.1 Behavioural approach

2.1.1 Invariance of shape

When subjected to ambiguous visual (or auditory) cues, observers (or listeners) typically experience perceptual reversals between alternative interpretations of the stimulus. In binocular rivalry, one of the most ancient and studied case of multistable perception, two different images (for instance, coloured gratings) are presented to one and the other eye. Rather than a fusion of both images, subjects experience a series of perceptual alternations between the image presented to one eye, and the other. Alternations are not regular, and form a sequence in which the elapsed time between reversal appears, to some extent, random. The properties of reversal times sequences for ambiguous displays have been studied intensively. The average time between subsequent reversals can differ widely between subjects and, for a given subject, can be strongly modulated by displays and stimulus qualities such as input levels [Fox and Herrmann, 1967, Borsellino et al., 1972, Walker, 1975, Zhou et al., 2004, Brascamp et al., 2005]. However, the shape of the corresponding reversal times distributions, does not: it remains astonishingly invariant between observers, regardless of displays, and stimulus qualities. In particular, it conserves a characteristic Gamma-like shape (with the coefficient of variation $c_v \approx 0.6$ and the skewness $\gamma_1 \approx 2c_v$). [Levelt, 1967, Blake et al., 1971, Walker, 1975, De Marco et al., 1977, Murata et al., 2003, Pastukhov and Braun, 2007]. This effect can be made particularly evident when reversal time densities are normalised to their average reversal duration (see **Fig.2.1**).

The precise shape of these distributions has been the subject of much debate and detailed investigations [De Marco et al., 1977, Lehky, 1988, Zhou et al., 2004, Brascamp et al., 2005, Shpiro et al., 2009], with several authors suggesting that other distributions, such as the log-normal distribution or the Weibull distribution, may provide better fitness to experimental observations than the traditionally used Gamma distribution. This study presents new evidences which may help settling this debate, by proposing a different way to probe the available data. Specifically,

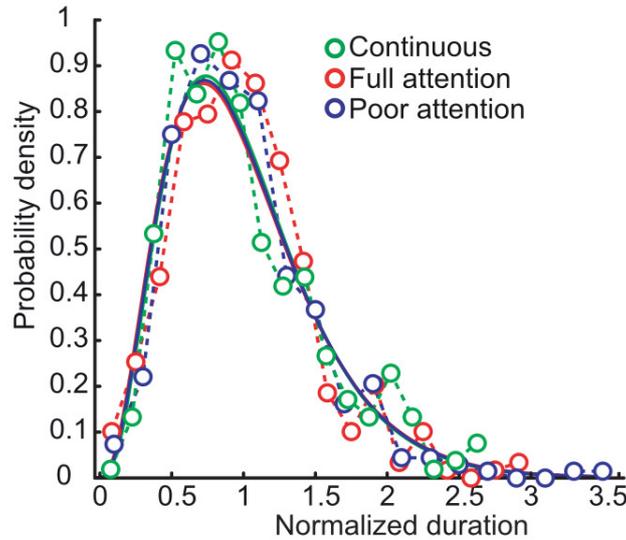


Figure 2.1: Illustration of the shape-invariance of reversal times distributions. Normalised distributions of dominance durations appear to conserve their characteristic shape, despite greatly different means. (*Green dotted lines*): continuous presentation, mean dominance time ≈ 4.9 s. (*Red dotted lines*): intermittent presentation with full attention, mean dominance time ≈ 25.9 s. (*Blue dotted lines*): intermittent presentation with poor attention, mean dominance time ≈ 37.4 s. Gamma distribution fits (*solid lines*) yielded: $c_v^{(cont)} = 0.51 \pm 0.1$, $c_v^{(full)} = 0.52 \pm 0.1$, and $c_v^{poor} = 0.51 \pm 0.1$. Reprinted from [Pastukhov and Braun, 2007] with the author’s permission.

we looked at the relationship between the mean of reversal times distributions and their higher-order *moments*. This is a simple but important nuance, which provides with a quantitative criterion other than the goodness-of-fit to identify which distribution may be most suited to account for the available data. In particular, it allowed us to disambiguate between two separate issues: the shape-invariance of reversal time densities and the actual shape of said densities. On the one hand, the constancy of distribution shape constitutes a ‘scaling property’, as it implies that higher moments scale as appropriate powers of the mean. On the other hand, the shape of the distribution imposes that higher moments take specific values in order to match behavioural observations. Throughout the course of this dissertation, we will justify that both issues should be tackled separately, and may indeed reflect different underlying dynamical features, as far as the neural processes driving conscious representations are concerned.

2.1.2 Invariance of moments

A large part of this work consists of comparing the moments' statistics of experimentally obtained sequences of perceptual reversals, with simulated sequences generated from a variety of models. A sequence of perceptual reversals obtained from an experimental trial or a computational model, can be defined by the sequence of random numbers $\{r_i\}$ corresponding to the absolute time at which each perceptual reversal has been reported by the subject. The purpose of the present study is to characterise the statistical properties of perceptual dominance durations, which are simply defined by the difference $t_i = r_{i+1} - r_i$ between successive reversal times. The sequence $\{t_i\}$ of dominance durations therefore corresponds to successive realisations t of a random variable T , with an associated probability density function (p.d.f.) P_T : the distribution of dominance durations, or equivalently, the reversal time density.

We can obtain the mean (noted μ_1), as well as higher-order central moments (noted μ_k , $k \geq 2$) of this distribution by computing ensemble averages (noted $\langle \cdot \rangle$) over all realisations of t .

$$\mu_1 \doteq \langle t \rangle = \int_0^{+\infty} t P_T(t) dt \quad (2.1)$$

$$\mu_k \doteq \langle (t - \mu_1)^k \rangle = \int_0^{+\infty} (t - \mu_1)^k P_T(t) dt \quad k \geq 2 \quad (2.2)$$

Rather than using central moments to characterise the shape of reversal times densities, we use normalised moments. At second order in fluctuations, we introduce the coefficient of variation c_v (CV):

$$c_v \doteq \frac{\mu_2^{1/2}}{\mu_1} \quad (2.3)$$

The CV of a given sequence of reversal times typically quantifies how regularly perceptual reversals occur. If $c_v \rightarrow 0$, we have a near constant alternation rate, meaning that times between subsequent reversals are almost the same and thus

alternation sequences would appear very regular. If $c_v \rightarrow 1$, times between subsequent reversals will be close to a Poisson process, and thus alternation sequences would appear very irregular. In the context of bistable perception, the CV of reversal sequences typically takes intermediate values ($c_v \approx 0.6$) [Fox and Herrmann, 1967, Walker, 1975, Pastukhov and Braun, 2007, Winkler et al., 2012]. The CV is a natural choice for studying the scaling property, because as average dominance periods may vary greatly, a dimensionless measure of fluctuations in the units of the mean is more desirable.

At third order in fluctuations, we introduce the skewness, noted γ_1 :

$$\gamma_1 \doteq \frac{\mu_3}{\mu_2^{3/2}} \quad (2.4)$$

Observed reversal time densities are noticeably asymmetric, and typically right-skewed, indicating that these densities have a right-tail longer than the left tail. However, the skewness is seldom reported or discussed in experimental and theoretical studies of multistable perception, perhaps because of the amount of data required to obtain reliable estimates. If the shape-invariance is indeed an emergent property of some underlying dynamical system, introducing this quantity may offer insights in some of its properties, which in turn may correspond to important functional features of corresponding processes at the neural level.

Formally, the scaling property requires the k -th central moment to scale with the mean as follows [Okamoto and Fukai, 2001]:

$$\frac{\sqrt[k]{\mu_k}}{\mu_1} = \text{constant} \quad k \geq 2 \quad (2.5)$$

At second and third order in fluctuations, this yields the following relations:

$$\mu_2 \propto \mu_1^2 \Rightarrow c_v = \text{constant} \quad (2.6)$$

$$\mu_3 \propto \mu_1^3 \Rightarrow \gamma_1 = \text{constant} \quad (2.7)$$

Furthermore, we will consider the ratio between the skewness and the CV rather

than the skewness alone. This is because, for a Gamma distribution, or an Inverse Gaussian distribution, two distributions which will appear recurrently throughout this thesis, the skewness is proportional to the CV ($\gamma_1 = 2c_v$ for the former, and $\gamma_1 = 3c_v$ for the latter). This rescaling of the skewness by the CV therefore allows direct comparison between observed distributions of dominance durations and a Gamma, or an inverse Gaussian distribution. In the following, this measure will be a tie-breaker between the different types of distributions which have been commonly used to fit behavioural data.

2.1.3 Experimental methods

All the relevant behavioural data presented in this study have been collected by Dr. Alexander Pastukhov, under the supervision of Pr. Jochen Braun (Institute for Cognitive Biology, Otto-von-Guericke Universität, Magdeburg), with the exception of data on auditory streaming, provided by I. Winkler and S. Denham [Bregman, 1994, Winkler et al., 2012]. Here, we provide some additional details on how experimental data relevant to this study have been collected. For the sake of consistency and generality, several different paradigms of multistable perception have been used.

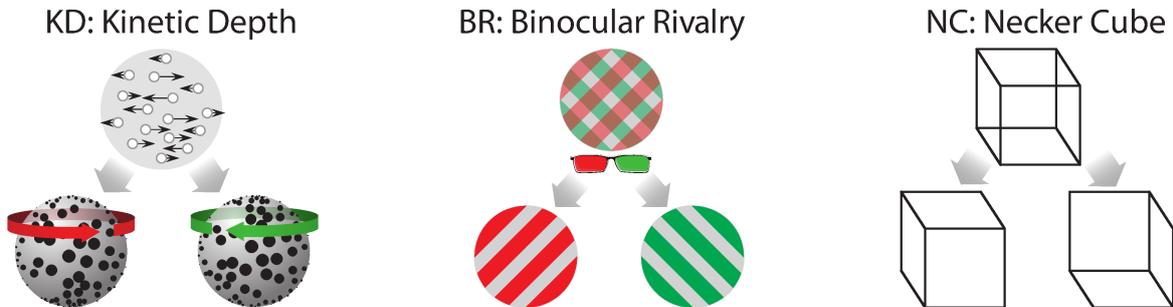


Figure 2.2: Examples of canonical multistable displays. (*Left*): In the kinetic-depth-effect (KDE), dots moving on a plane can be perceived as a rotating sphere. Leftward moving dots can be interpreted as forming the back surface of a leftward rotating sphere, or alternatively, as forming the front surface of a rightward rotating sphere. (*Middle*): In binocular rivalry, coloured grating with different orientations are presented to one and the other eye, perception then alternates between the two gratings (*Right*): The Necker cube, which in this particular orientation can be interpreted as having an outer vertex pointing left upward (left) or right downward (right).

Apparatus and General procedure

Stimuli were generated with MATLAB. Observers responded using a keyboard. Background luminance was kept at 36cd/m^2 . The experimental room was lit dimly (ambient luminance at 80cd/m^2). For visual multi-stability, we considered situations in which strong retinal inputs are spontaneously suppressed: ‘motion-induced-blindness’ or MIB, ‘binocular rivalry’ or BR [Campbell and Howell, 1972, Bonnef et al., 2001, Leopold and Logothetis, 1999], spontaneous reversals of illusory rotation in depth (‘kinetic depth effect’ or KDE) [Wallach and O’Connell, 1953, Sperling and Doshier, 1994], spontaneous shifts in the apparent direction of motion (‘moving plaids’, MP) [Adelson and Movshon, 1982, von Gruenau and Dube, 2005], and spontaneous changes in the three-dimensional appearance of line drawings (‘Necker cube’ or NC) [Meng and Tong, 2004]. The data sets on KDE, NC, and MP, as well as one of six data sets on BR, were published previously [Pastukhov and Braun, 2007, Pastukhov et al., 2013]. New data sets were collected on MIB and on BR (by Dr. Alexander Pastukhov, under the supervision of Pr. Jochen Braun). For auditory multi-stability, we considered spontaneous transitions between auditory segregation and integration (‘auditory streaming’, [Bregman, 1994, Winkler et al., 2012]). This data set was kindly provided by I. Winkler and S. Denham (personal communication). For choice reaction tasks, we re-analyzed published information on saccadic reaction times (‘saccade reaction time’, or SRT) [Carpenter, 2012] and on working memory retrieval times (‘memory reaction time’, or MRT) [Pearson et al., 2014]. Both studies report timing densities for individual observers.

Binocular rivalry:

Six observers participated in the experiment (4 male, 2 female). Stimuli were displayed on an LCD screen (EIZO ColorEdge CG303W, resolution 2560×1600 pixels, viewing distance was 104 cm, single pixel subtended 0.014° , refresh rate 60Hz) and were viewed through a mirror stereoscope. Chin and head rests were used to stabilize viewing position.

Two grayscale circular orthogonally-oriented gratings ($+45^\circ$ and -45°) were

presented foveally to each eye. Gratings had diameter of 1.6° , spatial period $2 \text{ cyc}/\text{deg}$. To avoid a sharp edge, grating contrast was modulated with Gaussian envelope (starting inner radius 0.6° , $\sigma = 0.2^\circ$). Tilt and phase of gratings was randomized for each block. Five contrast levels were used: 6.25%, 12.5%, 25%, 50%, and 100%. Contrast of each grating was systematically manipulated, so that each contrast pair was presented in two blocks (50 blocks in total). Each block was two minutes long and separated by a compulsory one-minute break. Observers reported on the tilt of the visible grating by continuously pressing one of two arrow keys. They were instructed to press only during exclusive visibility of one of the gratings, so that mixed percepts were indicated by neither key being pressed ($25 \pm 8\%$ of total presentation time). To facilitate binocular fusion, gratings were surrounded by a dichoptically presented square frame (outer size 9.8° , inner size 2.8°).

Motion-induced blindness:

Twenty observers participated in the experiment (12 male, 8 female). Stimuli were presented on a CRT screen (Iiyama VisionMaster Pro 514, iiyama.com, resolution 1600×1200 pixels, refresh rate 100 Hz). The viewing distance was 73 cm so that each pixel subtended approximately 0.019° . Target was a yellow circle (diameter 0.2°), presented 1° above the fixation. Mask rotated at 1 Hz and consisted of 8×8 grid of crosses (arm length 0.6° , inter-cross distance 0.15°). Observers reported episodes of target disappearance by keeping the space key pressed for the entire duration of the episode. Blocks lasted one minute and were separated by a compulsory 30 s. In total, 32 experimental blocks were measured for each observer.

2.1.4 Behavioural observations

Average dominance durations

Experimental observations are summarized in **Fig. 2.3**. For individual observers, mean dominance periods varied up to four-fold between the 25 contrast combinations (a 5×5 matrix of left and right eye contrasts), with mean dominance periods typically decreasing for increasing stimulus strength. Between different observers,

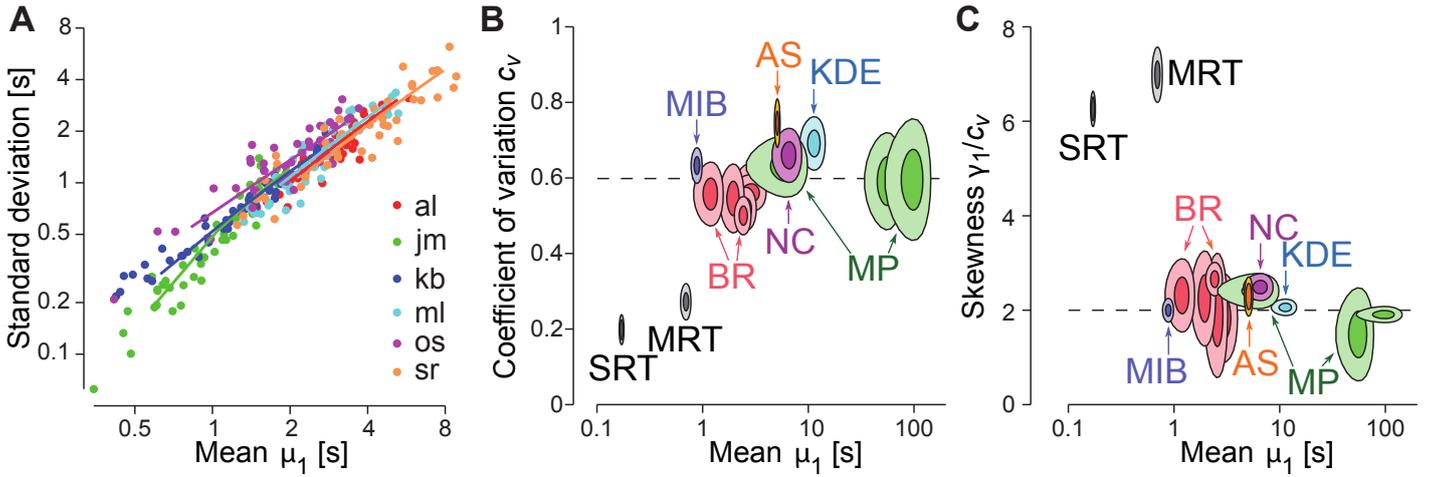


Figure 2.3: Illustration of the shape-invariance of reversal times distributions. **A** Standard deviation (SD) and mean μ_1 of the distribution of dominance durations. Individual values (round symbols) for 25 contrast combinations of binocular rivalry (BR) and 6 observers (different colors). Proportionality is maintained both across conditions (contrast combinations) and across observers. Linear regression (colored lines) yields a highly consistent $c_v = 0.66 \pm 0.04$ (mean \pm SD). **BC** Normalized moments for the distribution of dominance durations observed with different multistable perceptions: auditory streaming (AS), binocular rivalry (BR), kinetic depth effect (KDE), motion-induced blindness (MIB), moving plaids (MP), Necker cube (NC). Several stimulus conditions are shown for BR and MP (see text). For comparison, saccade reaction time (SRT) and memory reaction time (MRT) are included as well. **B** Coefficient of variation c_v , as a function of the mean μ_1 . All multistable situations exhibit comparable variability, with $c_v \approx 0.6$ (dashed line). **C** Skewness γ_1 , in multiples of c_v , as a function of the mean μ_1 . All multistable situations exhibit comparable skewness, with $\gamma \approx 2 c_v$ (dashed line). Shaded areas in **B** and **C** indicate one and two SEM across observers. Moreover, the distribution in question is considerably more variable (‘wider’) and less skewed (‘shorter-tailed’) than typical choice time distributions [Carpenter, 2012, Pearson et al., 2014]. *Reprinted from [Cao et al., 2016] with the co-authors’ permission.*

mean dominance periods differed by a further factor of four, with some participants experiencing greater switching rates than others. Between different display types, mean dominance periods ranged even more widely, as illustrated in **Fig. 2.3 BC**. Over six types of multistable visual displays (BR; kinetic depth effect, KDE; motion-induced blindness, MIB; moving plaids, MB; Necker cube, NC) and one multistable auditory scene (auditory streaming, AS), mean dominance periods varied over two orders of magnitude, from under 1 s to over 100 s.

The two choice reaction tasks illustrated here span the full range of human reaction times, from 100 *ms* to 1 *s* (**Fig. 2.3 BC**).

Higher-order statistics

Multistable phenomena maintain a remarkably strict ‘scaling property’ over different stimulus strengths, observers, and display types. The standard deviation of the dominance distribution is (nearly) proportional to the mean, as illustrated by the linear regression lines in **Fig. 2.3 A**. To assess the generality of this finding, this results was established for a variety of different visual and auditory conditions: AS, MIB, NC, and KDE, BR for stimulus contrasts $c \in \{0.06, 0.12, 0.25, 0.5, 0.6, 1.0\}$, and MP for continuous display, intermittent display, and intermittent display without attention [Pastukhov and Braun, 2007]. In each case, we established the mean, standard deviation (SD), CV, and skewness of the distribution of dominance periods. **Fig. 2.3 BC** illustrates the results in terms of the mean and SEM across observers. The CV remained consistently near $c_v \approx 0.6$ (**Fig. 2.3 B**) and the skewness remained consistently near $\gamma_1 \approx 2 c_v$ (**Fig. 2.3 C**). In other words, a scaling property was consistently maintained over all investigated situations, with values of the moments close to that of a Gamma distribution.

Note that the Log-normal or Inverse Gaussian are both significantly more skewed ($\gamma_1 \geq 3 c_v$). On the other hand, the Weibull distribution is significantly less skewed than a Gamma distribution for intermediate values of the CV ($\gamma_1 < 2 c_v$ if $c_v < 1$); a complete review of moment-ratio diagrams for these distributions can be found in [Vargo and Leemis, 2010].

$$\frac{\gamma_1}{c_v} = 2 \quad (\text{Gamma}) \quad (2.8)$$

$$\frac{\gamma_1}{c_v} = 3 \quad (\text{Inverse Gaussian}) \quad (2.9)$$

$$\frac{\gamma_1}{c_v} = 2 + e^{\sigma^2} \geq 3 \quad (\text{Log-normal}) \quad (2.10)$$

For comparison, we established the same statistics for choice reaction times.

In contrast to mean dominance periods, choice reaction times have very different statistical properties. Firstly, a two-fold range of distribution means with stimulus strengths or between observers appears typical [Carpenter, 2012, Pearson et al., 2014]. This is not nearly as variable as what is observed for typical reversal times. Secondly, they distribute very differently, with distribution shapes combining smaller variance ($c_v \approx 0.2$, i.e. sharper than in the multistable case) with greater skewness ($\gamma_1 > 6c_v$, i.e. with heavier tails than in the multistable case). If multistability can be thought of as a limiting case of a discrimination task, where competing percepts can not be disambiguated, such disparity in how the timing of perceptual events are distributed would indicate that multistable displays cause sensory evidences to be accumulated and evaluated in a radically different manner. In that respect, multistable perception could be more than a fringe effect, artificially crafted in laboratories. Understanding its peculiar statistics may reveal a more general mechanism and provide strong constrains on the properties of the underlying neural substrates.

The available behavioural data on multistable perception define a set of constrains which can be summarised as follows:

- (i) *Wide behavioural range*: Distribution mean varies widely across conditions and decreases with stimulus strength.
- (ii) *Scaling property*: Distribution shape is highly preserved across conditions (stimulus strengths, observers, and display types)
- (iii) *Distribution shape*: Distribution variance is large ($c_v \approx 0.6$) while the distribution skewness is small $\gamma_1 \approx 2c_v$, in comparison to choice reaction times, for instance.

The characterisation of the scaling property via its moments constitutes the basis around which this study articulates. Throughout this work, results obtained from each computational model will be systematically confronted to these experimental evidences. In particular, **Fig. 2.3 BC** will be reproduced by computational models and referred to on numerous occasions.

2.2 Modelling approach

2.2.1 First-passage-time framework

The timing of perceptual events, such as reaction times or reversal times, is thought to reflect the gradual, noisy accumulation of sensory information, such that a perceptual decision, or reversal, occurs when sufficient amount of evidences has been collected. A commonly used simplification is to reduce the dynamics cortical populations performing this accumulation to the stochastic dynamics of a single stochastic decision variable, corresponding to the level of activity in neural populations driving the decision. Stimulus onset causes this activity to accumulate, until a set threshold is reached and a decision ensues: this is known as a first-passage-time (FPT) framework. While this approach may seem a crude simplification, such framework provides a convenient and concise terms to model the timing of perceptual events, and has proved successful in accounting for several aspects of reaction times distributions in the context of motion-discrimination tasks [Ratcliff and Smith, 2004, Smith and Ratcliff, 2009]. For a number of stochastic processes, the FPT problem is mathematically tractable, providing a direct quantitative link between neural dynamics and behaviour, and between models' properties and experimental constrains.

Our study sought to extend this diffusion-to-bound framework to the timing of perceptual reversals. Our first objective was to compare the ability of various stochastic processes to reproduce the scaling property of multistable perception. Investigating the qualitative reasons behind their failure, or success to do so allowed us to single-out the essential dynamical features from which the scaling property may originate. As is the case in 2AFC tasks, multistable phenomena may involve at least two decision variables (one corresponding to each perceptual appearance). For instance, a given perceptual appearance would take dominance when its associated decision variable gains a sufficient advantage over its concurrent (evidence-biased competition). Presumably, following a reversal, these roles would

be switched, so that the decision variable associated with the newly dominant appearance decreases, while the decision variable associated with the newly suppressed appearance increases, until a differential threshold is reached again and the cycle starts anew.

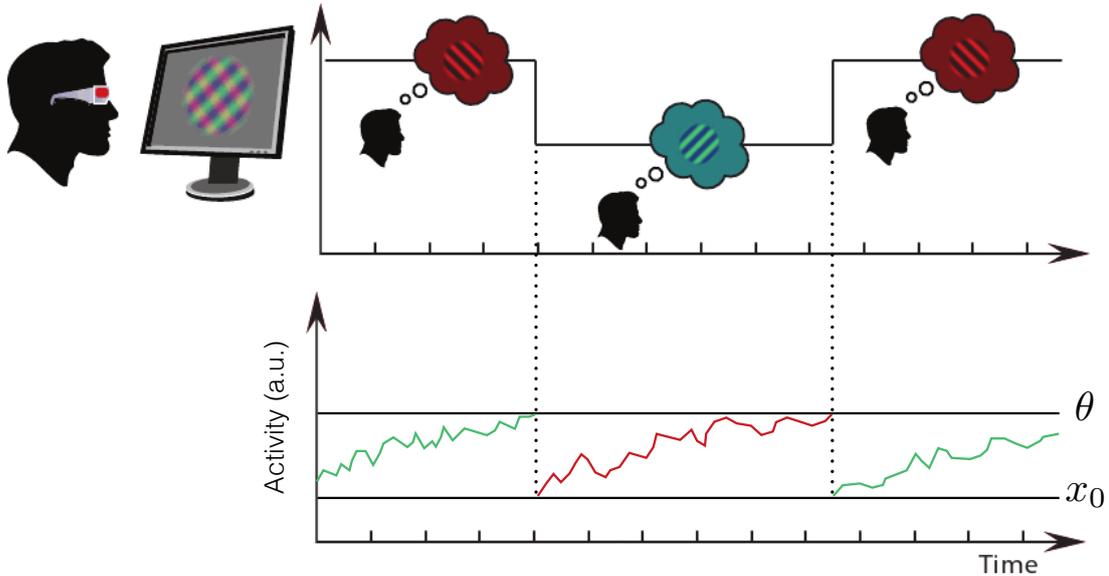


Figure 2.4: Illustration of the correspondence between first-passage-times and reversal times during the presentation of an ambiguous display. Dominance periods for a given perceptual appearance (e.g. the green gratings) correspond to the time-to-threshold (from x_0 to θ) of the decision variable associated with the currently suppressed perceptual appearance (e.g. sensory evidences supporting the red grating), and conversely. Reversals are initiated when the activity reaches a threshold, as marked by vertical dotted lines.

We first consider a simplified version of this framework, in which we are not concerned with implementing the details of dominance reversals; this will be the object of Chapter 5. We initially assume that the conclusion of one stochastic accumulation coincides with the start of another such accumulation (see **Fig. 2.4**). In order to study the scaling property, we only need to consider the properties of the FPT density under the variation of a *single* input-level: this is enough to obtain dominance periods matching the range of observed behaviour. Therefore, it is also sufficient to consider the FPT statistics of only one of the two percepts,

simply because we do not yet need to consider the joint-dependence of reversal times on both input-levels (Levelt's propositions). This will also be treated in Chapter 5, where we will relax this assumption and detail the mechanisms driving percept reversals, so that the subtleties of Levelt's propositions can be addressed.

To summarize: in an initially reduced description, we consider the noisy accumulation of activity, noted $x(t)$, over time, in a neural population, reflecting the accumulation of supporting evidences in favour of a given perceptual representation. Assuming a time t_0 directly following a perceptual reversal, so that the corresponding percept is currently suppressed, and the activity in the supporting population is at a low value of activity $x(t_0) = x_0$. We assume $x(t)$ to gradually accumulate, in a random-walk fashion, until it reaches a perceptual threshold θ , which we assume triggers the next perceptual reversal, so that the associated percept now takes dominance. This means that the time-to-threshold, or equivalently FPT, from x_0 to θ , noted $T_{x_0\theta}$, of the decision variable supporting the suppressed representation, corresponds to the dominance duration of the competing (then dominant) representation. We will then examine the moments of the associated FPT distribution, and whether they obey the scaling property under variations of stimulus strength.

2.2.2 Comparison criterion

In treating multistable phenomena as a first-passage time problem, we assume that the stochastic intervals between spontaneous reversals characterize an underlying random walk performed by microscopic states and reflected in collective accumulation of activity. A large part of this Chapter consists of drawing both qualitative and quantitative comparisons between the properties of first-passage-time densities (or equivalently reversal time densities) for different random walk models. In general, the dynamics of random walks, such as drift-diffusion models, consist of a *deterministic component*, reflecting to the average dynamics over many realisations of the process, and a *stochastic component*, reflecting to the trial-to-trial variability of the process. The time-to-threshold properties of different stochastic processes can therefore be compared by studying the rate at which they accumulate activity over

time, driven by the deterministic component, and the rate at which they disperse activity over time, driven by the stochastic component. For a given random walk $x(t)$, the accumulation and dispersion rate can be defined as the time-derivative of the infinitesimal mean $\langle x(t) \rangle$ and infinitesimal variance $\sigma_x(t)^2 = \langle (x(t) - \langle x(t) \rangle)^2 \rangle$, respectively (here $\langle \cdot \rangle$ denotes the ensemble average over multiple realizations of the stochastic process $x(t)$):

$$\nu_{drift} \doteq \frac{d}{dt} \langle x(t) \rangle \quad (2.11)$$

$$\nu_{noise} \doteq \frac{d}{dt} \langle (x(t) - \langle x(t) \rangle)^2 \rangle \quad (2.12)$$

Depending on the type of diffusion process under study, such rates may be input-dependent, meaning they may be modulated by the stimulus strength s , but also state-dependent and change over time with the instantaneous activity $x(t)$:

$$\nu_{drift} = \nu_{drift}(x, s), \quad \nu_{noise} = \nu_{noise}(x, s) \quad (2.13)$$

In particular, we expect that ν_{drift} will increase for increasing stimulus strength, so that the activity would accumulate at greater rate for stronger stimuli. This implies that the mean FPT $\mu_1 = \langle t \rangle$ is expected to decrease with input s , consistently with experimental observations. The same is true for the central moments μ_2 and μ_3 , as previously highlighted in **Fig. 2.3 BC**.

It is intuitive that FPT density depends sensitively on deterministic forces, stochastic forces and on the activity range traversed. In particular, its shape will depend on the *balance between deterministic and stochastic forces*. A change to any single variable, deterministic force, stochastic force, or traversed range, will alter the FPT density and all of its moments. For instance, changing the deterministic force, by increasing or decreasing the input-level, will affect the accumulation rate of sensory evidences, causing the distribution to shift its mean. Importantly, the balance between deterministic and stochastic components of the dynamics will

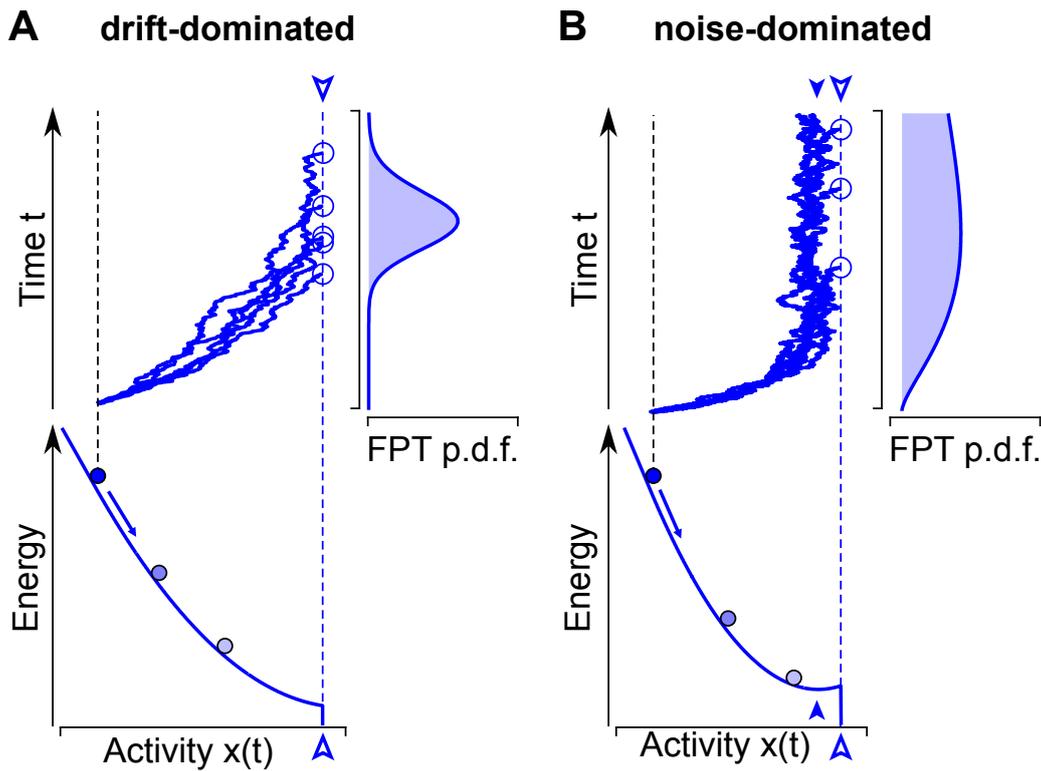


Figure 2.5: First-passage times (FPTs) of a threshold level θ by stochastic neuronal activity $x(t)$. In this framework, a perceptual decision or reversal is triggered when $x(t) \geq \theta$ for the first time. **A** In a ‘drift-dominated regime’ (DDR), deterministic forces drive activity to θ (open arrow) and beyond. **B** In a ‘noise-dominated regime’ (NDR), deterministic forces drive activity merely to a steady-state (filled arrow), some distance below θ (open arrow). Thus, θ may be reached only with the help of stochastic forces (noise). **AB:** In both regimes, individual realizations of neural activity $x(t)$ develop from an initial level (black dashed line) to a threshold level (blue dashed line, open arrows). Due to stochastic factors, every realization reaches threshold at a different time (open circles). This variability results in a probability distribution of FPT, which is illustrated on the right. Deterministic forces may be visualized in terms of an energy landscape, which is shown below. The sign of the energy gradient at threshold (open arrow) distinguishes drift- and in NDRs. *Reprinted from [Cao et al., 2016] with the co-authors’ permission.*

also be sensitive to variations of any of these parameters. If this balance shifts in favour of the former or the latter, so that threshold crossings become more regular, or more irregular, we can expect the distribution shape to sharpen, or widen, respectively. In general, first and higher moments will not change as to verify a scaling property, as this would require changes in of the higher moments to scale with changes of the mean.

Therefore, one important point to consider is the manner in which a given random walk model approaches and crosses a given threshold. Naturally, if a specific distribution shape is to be obtained *and* maintained for a wide range of stimulus strength, this should constrain the dynamics toward the reversal threshold. The aim of this chapter is precisely to identify these constraints.

Some of the diffusion models under study may operate in two distinct regimes, illustrated in **Fig. 2.5**. In a ‘drift-dominated regime’ (DDR), deterministic forces dominate, and are sufficient to drive activity over the threshold. Stochastic forces merely introduce some variability, reflected in the spread and overall shape of the distribution. In a ‘noise-dominated regime’ (NDR), deterministic forces drive activity toward a steady-state, which lies some distance below threshold. Stochastic forces are needed to drive activity away from steady-state (filled arrow), and over the threshold (open arrow). The difference is best appreciated in an effective energy landscape: the energy gradient takes opposite sign at threshold (**Fig. 2.5** open arrows).

In a DDR, the threshold is located below the asymptotic value ($x_{in} > \theta$): this will yield typical average FPT smaller than their characteristic time-constant τ (so that $\langle t \rangle < \tau$). In a NDR, the threshold is located below the asymptotic value ($x_{in} < \theta$): this will yield typical average FPT greater than their characteristic time-constant τ (so that $\langle t \rangle > \tau$). This is because the typical relaxation time to the asymptotic value effectively introduces a refractory period: the system must first reach its asymptote, which it does in multiples of τ . Then, it must cross the threshold with the aid of fluctuations around equilibrium.

Importantly, reversal times and densities produced in a drift- or noise-dominated regime will be affected by changes in input-levels in a significantly different manner. In the DDR, the threshold-crossing is ensured by the preponderance of the deterministic component of the dynamics over its stochastic components. The average FPT will mostly depend on deterministic factors while the distribution shape will depend on the balance between deterministic and stochastic factors. This bal-

ance, and more specifically, whether or not it is maintained under modulation of the stimulus strength, is a key element in replicating the statistical properties of multistable dynamics.

Note that in the NDR, there is no such balance: threshold-crossing exclusively will almost exclusively rely on the stochastic component of the dynamics, and specifically in the amplitude of fluctuations around the steady-state. Therefore we can already expect qualitative differences in how distribution shapes are modulated with the stimulus strength, depending on whether the accumulation is performed in either regime. As we will see in later chapters, several features of multistable dynamics strongly rely on the accumulation of sensory evidences operating in a DDR. Also, in the limit of strong DDRs, i.e. when the threshold θ and the activity $x(t)$, also, remain far below the input-dependent asymptotic value $x_{in}(s)$, accumulation and dispersion rates are independent of $x(t)$ and become stationary:

$$\nu_{drift}(x, s) \xrightarrow{x \ll x_{in}} \nu_{drift}^*(s) \quad (2.14)$$

$$\nu_{noise}(x, s) \xrightarrow{x \ll x_{in}} \nu_{noise}^*(s) \quad (2.15)$$

In this case, the stationary drift and dispersion rates ν_{drift}^* and ν_{noise}^* , and the threshold θ satisfy a simple relation, allowing for the required balance between deterministic forces, stochastic forces, and activity range (or distance to threshold) to be estimated. This provides a minimal criterion to compare how well different diffusion processes may, a priori, satisfy the scaling property. Moreover, we can explicit the correct interdependence between drift and noise terms, for which the required balance between deterministic and stochastic forces can hold under large stimulus strength modulations. We will see that, for a class of model which initially could not produce scale-invariant FPT distributions, the necessary corrections can be derived, and subsequently re-injected so that the scaling property is obtained.

2.2.3 Sufficient condition for the scaling property

For any drift-dominated accumulation (i.e. so that $x \ll x_{in}$), the variance of FPTs μ_2 is proportional to variance of activity σ_x^2 at time $\langle t \rangle$, so that the respective

coefficients of variation are approximately the same:

$$\sqrt{\mu_2} \propto \sigma_x \quad \Rightarrow \quad c_v^{(t)} \approx c_v^{(x)} \quad (2.16)$$

Also, the drift and dispersion rates can be simply expressed as:

$$\begin{aligned} \nu_{drift}^* \langle t \rangle &= \theta \\ \nu_{noise}^* \langle t \rangle &= \sigma_x^2 \end{aligned} \quad (2.17)$$

Using (2.16) and (2.17), we obtain an expression for the CV of dominance durations as a function of the drift and dispersion rates:

$$c_v^{(t)} = \frac{\sigma_x}{\theta} = \sqrt{\nu_{noise}^* / \theta \nu_{drift}^*} \quad (2.18)$$

Experimental evidence shows that increased input *decreases* the mean FPT $\langle t \rangle$ through the threshold θ *decreases*. It follows that increased input must *increase* accumulation rates ν_{drift}^* , so that:

$$\langle t \rangle = \theta / \nu_{drift}^*, \quad \nu_{drift}^* = f(s), \quad \partial f / \partial s > 0 \quad (2.19)$$

If ν_{noise}^* remains *constant* with input (this is typically the case when a constant noise parameter is assumed), it further follows that the variance at $\langle t \rangle$ decreases with ν_{drift}^* . This can be qualitatively understood from **Fig. 2.6**. As ν_{drift}^* increases, $\langle t \rangle$ decreases, leaving less time for the distribution of $x(t)$ to spread, so that the variance of $x(t)$ at time $\langle t \rangle$ (σ_x^2), also decreases.

$$\sigma_x^2 \propto f^{-1}(s) \quad (2.20)$$

Therefore, from (2.18), we expect the CV to decrease with the square root of

ν_{drift}^* , so that no scaling property is obtained:

$$c_v \propto f^{-1/2}(s) \quad (2.21)$$

However, if both ν_{noise}^* and ν_{drift}^* *increase* with input in the same proportions, the CV remains constant and the necessary condition for a scaling property is satisfied.

$$\nu_{noise}^* \propto \nu_{drift}^* \propto f(s) \quad \Rightarrow \quad c_v \propto \sqrt{\nu_{noise}^* / \theta \nu_{drift}^*} = const \quad (2.22)$$

Note that for simplicity, we have established the sufficient condition for the scaling property at second-order only. We will see in what follows that, in practice, in all cases studied, this will be enough to provide sufficient qualitative understanding. In principle, similar relations would be required for the skewness, and other moments, which would presumably rely on higher-order equivalents of expressions 2.11 and 2.12.

In what follows we will see that, in the case of Gaussian diffusion models, the required dependence of drift- and dispersion-rates to satisfy the scaling property correspond to a modulation of the unit of time with input-levels. This naturally changes the average first-passage-time without affecting the shape of the FPT distribution.

2.3 Gaussian diffusion-to-bound

In this section we compare the FPT of several continuous random walks $x(t)$ whose distribution obey the following Fokker-Planck equation (Gaussian drift-diffusion processes, see [Risken, 1984]):

$$\tau \frac{\partial p(x, t)}{\partial t} = -\frac{\partial}{\partial x} (\mu(x, x_{in}) p(x, t)) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (\sigma^2(x, x_{in}) p(x, t)) \quad (2.23)$$

where $\mu(x, x_{in})$ and $\sigma(x, x_{in})$ represent ‘accumulation’ or ‘drift’ and ‘dispersion’ or ‘noise amplitude’, respectively. In general, both terms may depend on the current state x and on external input x_{in} (we restrict ourselves to drift and diffusion

terms without any explicit dependence on time). Equation (2.23) describes the temporal evolution of density $p(x, t)$ over time under the influence of the drift and diffusion coefficients. The drift can be thought as the *deterministic component* of the dynamics, while the diffusion can be thought as the *stochastic component* of the dynamics. The former tends to shift the density $p(x, t)$ over time (for a time interval dt , the shift will be $\sim \mu(x, x_{in})dt$), without affecting its shape, while the latter tends to widen it over time (for a time interval dt , the width will increase by $\sim \sigma(x, x_{in})\sqrt{dt}$), as illustrated in **Fig. 2.6**, in a simpler case where drift and diffusion terms are constant (also known as the Wiener process with drift).

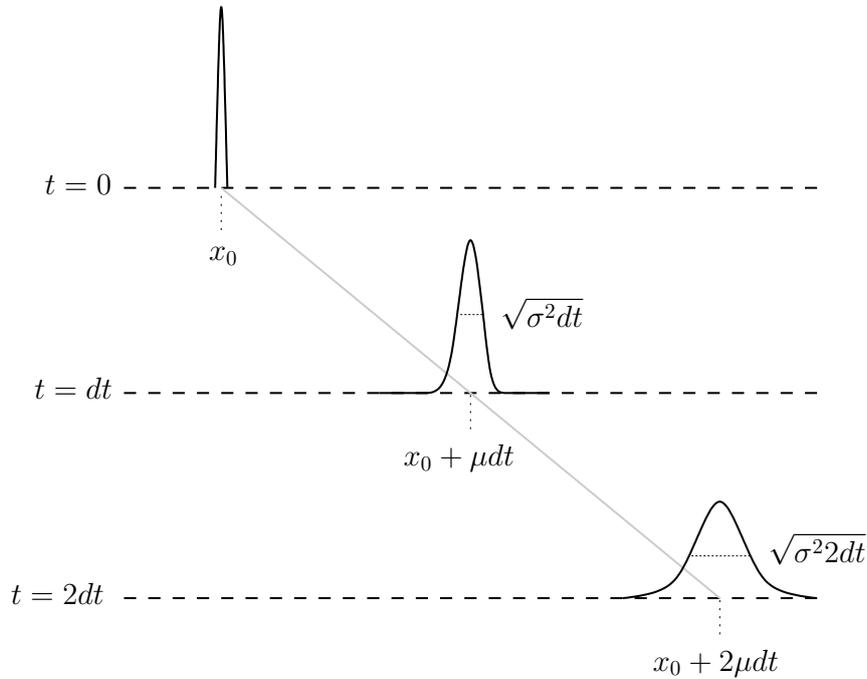


Figure 2.6: Illustration of the effect of the drift and diffusion terms on the evolution of a probability density $p(x, t)$, initially set to $p(x, 0) = \delta(x - x_0)$. In this example, $\mu(x, x_{in}) = \mu$ and $\sigma(x, x_{in}) = \sigma$ (Wiener process with drift).

An equivalent description of equation (2.23) is known as the Langevin equation. Rather than the evolution of the probability density of a random variable x , the Langevin equation describes the temporal evolution of a single realisation of random variable x . If one were to generate multiple realizations of x over time, the density

$p(x, t)$ would be retrieved. The stochastic differential equation for x is as follows:

$$\tau \frac{dx}{dt} = \mu(x, x_{in}) + \sqrt{\tau} \sigma(x, x_{in}) \xi(t) \quad (2.24)$$

where $\xi(t)$ is a unitary Gaussian ‘white noise’ (i.e. distributed as $\mathcal{N}(0, 1)$ where \mathcal{N} is the normal distribution), defined as:

$$\langle \xi(t) \rangle = 0, \quad \langle \xi(t) \xi(t') \rangle = \delta(t - t') \quad (2.25)$$

Under this formulation, it is clear that the drift term accounts the deterministic component of the dynamics, related to the average realisation over multiple trials, while the diffusion term accounts for the stochastic component of the dynamics, which differs from trial to trial. If we average out the stochastic component, we are left with the drift alone. Perceptual decisions are commonly modelled in terms of diffusion-like processes [Smith, 2000, Ratcliff and Smith, 2004, Smith and Ratcliff, 2009, Smith et al., 2014].

We start our analysis with two well-known and documented processes. The first one is the ‘Wiener process with drift’ (WP), which accumulates linearly, and the second one, the ‘Ornstein-Uhlenbeck process’ (OU), for which accumulation varies with activity and vanishes for at equilibrium. Whereas mean accumulation (drift) is assumed to vary with input, stochastic dispersion (noise) are usually kept constant. Next, we will show that a ‘scaling property’ can not be obtained if only accumulation rate, but not dispersion rate, changes with input. We then modify these processes, to consider scenarios where accumulation and dispersion change proportionally with sensory input s (through $x_{in}(s)$), so that they obey the sufficient condition for the scaling property derived in (2.22).

2.3.1 Standard Wiener process

We first study the scaling property in a WP with drift described by the following Langevin equation:

$$\tau \frac{dx}{dt} = x_{in} + \sqrt{\tau} \sigma \xi(t) \quad (2.26)$$

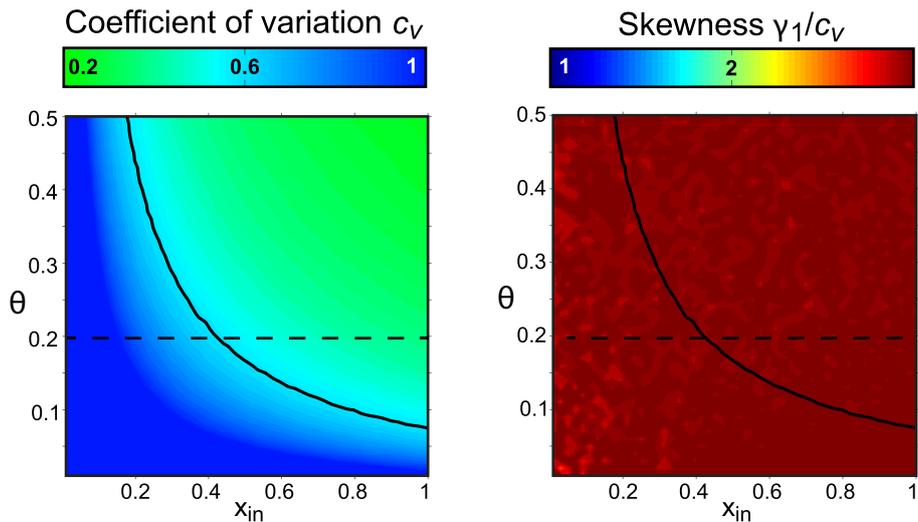


Figure 2.7: WP with input-dependent drift and constant noise: normalized moments of FPT density as functions of input x_{in} and threshold θ , (both in units of σ). (Left): coefficient of variation c_v . (Right): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$. Dashed line marks a particular choice of threshold θ . Reprinted from [Cao et al., 2016] with the co-authors' permission.

Here, the drift term $\mu(x, x_{in}) = x_{in}$ is constant and non-vanishing, corresponding to a linear integrator dynamics. The noise term σ is set to a constant. The FPT distribution for this process is known to be an inverse Gaussian (or Wald) distribution (see [Tuckwell, 1988/2008]). Assuming the initial condition $x_0 = 0$, the moments write:

$$\mu_1 = \frac{\theta \tau}{x_{in}} \quad (2.27)$$

$$c_v = \frac{\sigma}{\sqrt{x_{in} \theta}} \quad (2.28)$$

$$\gamma_1 = \frac{3\sigma}{\sqrt{x_{in} \theta}} = 3 c_v \quad (2.29)$$

In order to ensure a threshold crossing within a finite time, the drift term x_{in} must be strictly positive. In these conditions, this means that the escape through the absorbing barrier at threshold θ is always drift-dominated. Moments of the FPT density as a function of the input level and the value of the threshold are illustrated in **Fig. 2.7**.

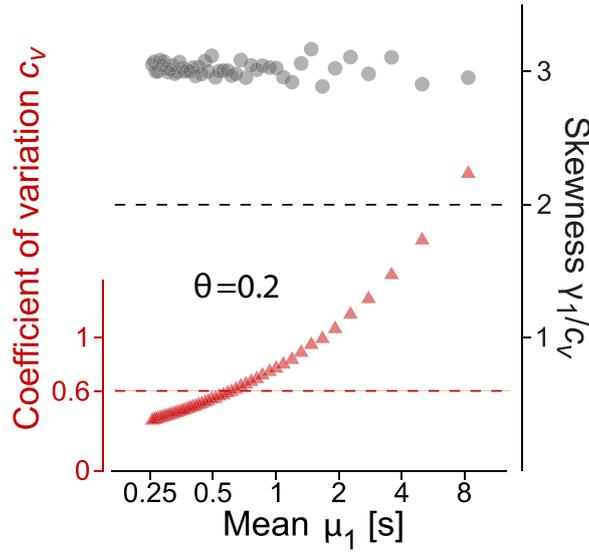


Figure 2.8: WP with input-dependent drift and constant noise: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). For constant θ (value given by inset), μ_1 decreases as x_{in} increases. The scalar property is not satisfied for the CV, and the ratio γ_1/c_v 3 is too large. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

Whereas the mean FPT μ_1 decreases with x_{in} and increases with threshold θ , the coefficient of variation c_v decreases with both input x_{in} and θ . The skewness is consistently high with $\gamma_1 = 3c_v$. The experimentally observed ‘scaling property’ is not reproduced (**Fig. 2.8**).

Alternatively, we can obtain the same formula for the CV as in (2.28) by computing the drift and diffusion rates for the WP using (2.11) and (2.12) and re-injecting the obtained expressions in (2.18). The mean and variance of the WP are well-known:

$$\langle x(t) \rangle = \frac{x_{in}t}{\tau} \quad (2.30)$$

$$\langle x(t)^2 \rangle - \langle x(t) \rangle^2 = \frac{\sigma^2 t}{\tau} \quad (2.31)$$

Therefore:

$$\nu_{drift} = \frac{x_{in}}{\tau} \quad (2.32)$$

$$\nu_{noise} = \frac{\sigma^2}{\tau} \Rightarrow c_v = \frac{\sigma}{\sqrt{x_{in}\theta}} \quad (2.33)$$

2.3.2 Modified Wiener Process

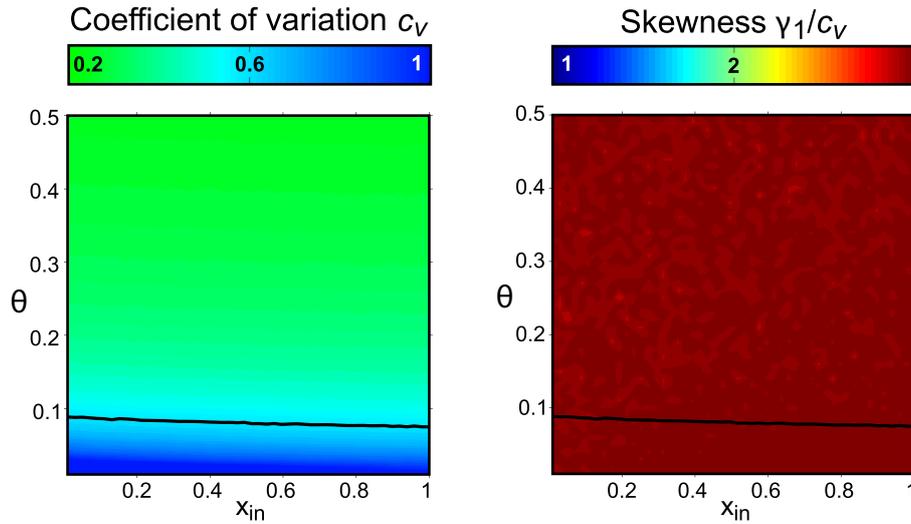


Figure 2.9: WP with input-dependent drift and noise: normalized moments of FPT density as functions of input x_{in} and threshold θ (both in units of σ). (*Left*): coefficient of variation c_v . (*Right*): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$. Dashed line marks a particular choice of threshold θ . *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

We now modify the constant- σ WP so that a scaling property is satisfied, and (2.22) holds. The correct mapping between noise and input is:

$$\sigma^2 = \frac{x_{in}}{\alpha} \quad (2.34)$$

So that the CV becomes:

$$c_v = \frac{1}{\sqrt{\alpha\theta}} \quad (2.35)$$

For a given threshold θ , constant α can be adjusted for the value of c_v to match with behavioural observations, i.e. $c_v = 0.6$. Moments of the FPT density as a function of the input level and the value of the threshold are illustrated in **Fig. 2.9**.

This time, as expected, the normalized moments c_v and γ_1 remain constant with x_{in} , while the mean FPT spans a wide range of values, as illustrated in **Fig. 2.10**.

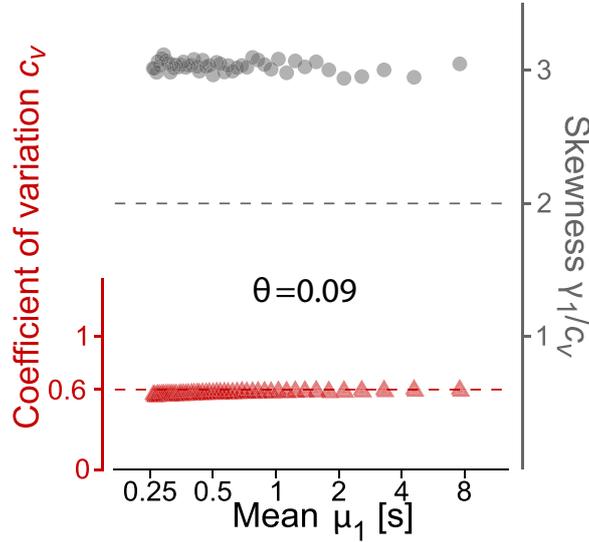


Figure 2.10: WP with input-dependent drift and noise: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). For constant θ (value given by inset), μ_1 decreases as x_{in} increases. The scalar property is satisfied for both the CV and the Skewness, but the ratio $\gamma_1/c_v = 3$ is too large. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

The qualitative explanation is straightforward: in the previous case with constant σ , increasing the input strength (or equivalently, the drift term x_{in}) yields an always steeper accumulation of activity, which in turns produces more and more regular threshold-crossings. This is because for increasing input-levels, the balance between the deterministic and stochastic components of the dynamics gradually moves in favour the deterministic component, causing the CV to decrease beyond the behavioural range. Our modification sets the amplitude of the noise terms to increase as x_{in} increases. This effectively restores the required balance, adaptively, for different stimulus intensity, so that a wide range of mean FPT can be accessed without breaking this balance. The skewness however, is too large to match behavioural observations $\gamma_1/c_v = 3$, which corresponds to an inverse Gaussian distribution.

2.3.3 Standard Ornstein-Uhlenbeck Process

We now study the scaling property in another widely used random walk model described by the following Langevin equation:

$$\tau \frac{dx}{dt} = x_{in} - x + \sqrt{\tau} \sigma \xi(t) \quad (2.36)$$

This process is known as the OU process. In this case, the drift term $\mu(x, x_{in}) = x_{in} - x$ depends, as in the WP, on the input strength, but is also state-dependent. In particular, it depends negatively on the activity $x(t)$, and vanishes when the system reaches equilibrium, i.e. when $x(t) = x_{in}$. For $\theta > x_{in}$, the drift vanishes before the threshold is reached, and the system becomes noise-driven.

In both drift- and noise-dominated regimens, closed-form expressions for the moments of the FPT density exist, in terms of infinite series [Inoue et al., 1995], or nested integrals [Brunel, 2000]. Moments and their respective dependence on the input strength x_{in} and the threshold value θ are illustrated in (Fig. 2.11).

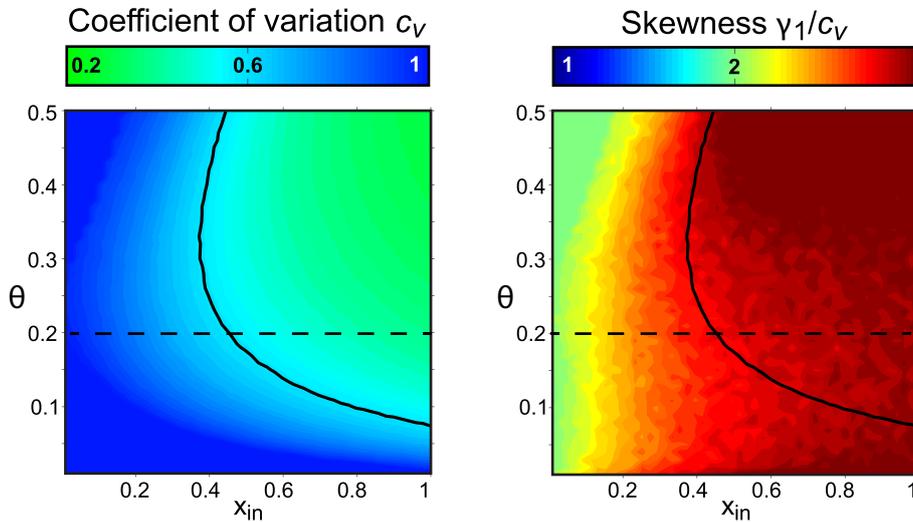


Figure 2.11: OU process with state- and input-dependent drift and constant noise: normalized moments of FPT density as functions of input x_{in} and threshold θ (both in units of σ). (Left): coefficient of variation c_v . (Right): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$. Dashed line marks a particular choice of threshold θ . Reprinted from [Cao et al., 2016] with the co-authors' permission.

As in the case of the WP, the coefficient of variation c_v varies with x_{in} and

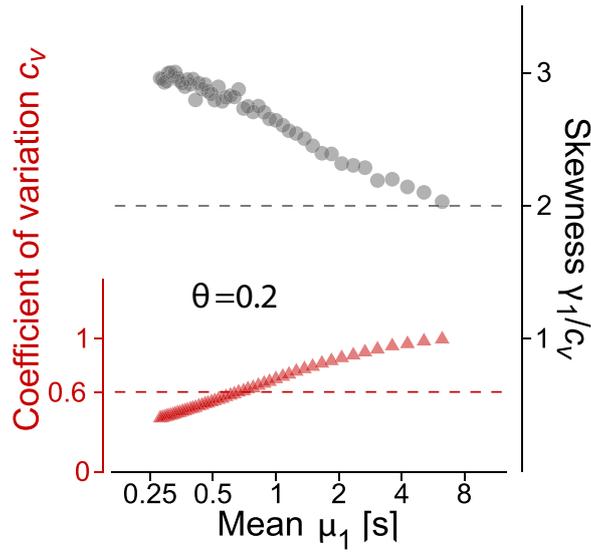


Figure 2.12: OU process with input-dependent drift and constant noise: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). For constant θ (value given by inset), μ_1 decreases as x_{in} increases. The scalar property is not satisfied for either the CV, or the Skewness. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

the skewness $\gamma_1 \lesssim 3c_v$ is too high (**Fig. 2.12**): the ‘scaling property’ can not be obtained.

Computing the drift and diffusion rates for the OU process leads to the same conclusion as for the WP with constant noise. Mean and variance are also known:

$$\langle x(t) \rangle = x_{in}(1 - e^{-\frac{t}{\tau}}) \quad (2.37)$$

$$\langle x(t)^2 \rangle - \langle x(t) \rangle^2 = \frac{\sigma^2}{2}(1 - e^{-\frac{2t}{\tau}}) \quad (2.38)$$

We obtain the drift and dispersion rates:

$$\nu_{drift} = \frac{x_{in}}{\tau} e^{-\frac{t}{\tau}} \quad (2.39)$$

$$\nu_{noise} = \frac{\sigma^2}{\tau} e^{-\frac{2t}{\tau}} \quad (2.40)$$

And in the limit of a DDR, we obtain:

$$\nu_{drift}^* = \frac{x_{in}}{\tau} \quad (2.41)$$

$$\nu_{noise}^* = \frac{\sigma^2}{\tau} \Rightarrow c_v = \frac{\sigma}{\sqrt{x_{in}\theta}} \quad (2.42)$$

2.3.4 Modified Ornstein-Uhlenbeck Process

The same linear dependence (2.34) of the infinitesimal variance on stimulus-level can be used to insure the scaling property of the OU process.

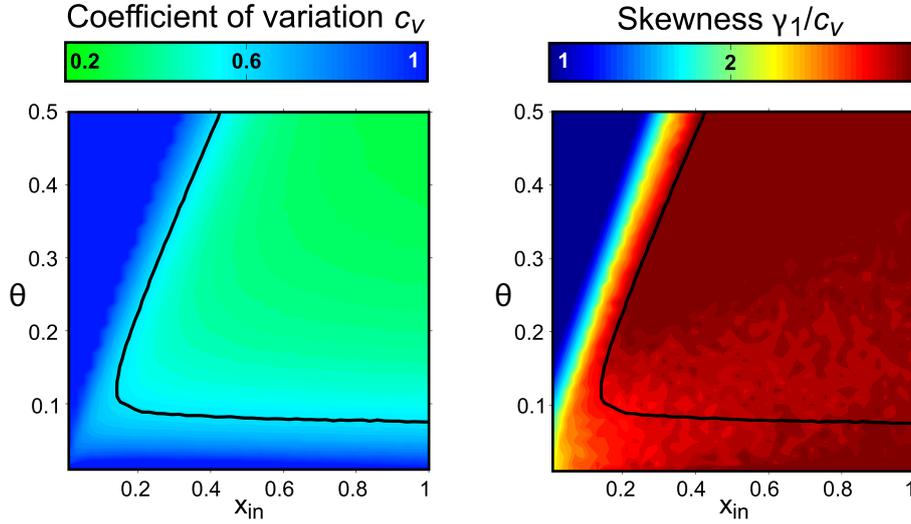


Figure 2.13: OU process with input-dependent drift and noise: normalized moments of FPT density as functions of input x_{in} and threshold θ (both in units of σ). (Left): coefficient of variation c_v . (Right): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$. Dashed line marks a particular choice of threshold θ . Reprinted from [Cao et al., 2016] with the co-authors' permission.

Moments of the FPT density are illustrated in **Fig. 2.13**. The two separate regimes are clearly visible: a regime where the threshold remains low and constant and a regime where the threshold increases linearly with input strength. In the limit of strong DDRs, we retrieve the FPT density of the modified WP (inverse Gaussian).

As before, skewness $\gamma_1 \approx 3 c_v$ is consistently higher than experimentally observed (**Fig. 2.14**, right).

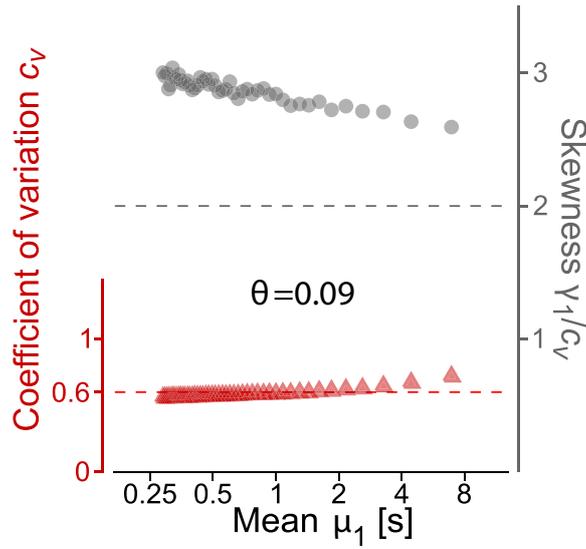


Figure 2.14: OU process with input-dependent drift and noise: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). The scalar property is satisfied for both the CV and the Skewness, but the ratio γ_1/c_v 3 is too large. Reprinted from [Cao et al., 2016] with the co-authors' permission.

2.4 Summary

To summarize, maintaining the proportionality $\nu_{\text{noise}}^* \propto \nu_{\text{drift}}^*$ ensures the constancy of normalized moment in a DDR. When the threshold θ is held constant, this requires the noise amplitude σ to become input-dependent. As explained by (2.18), the variability in the FPT statistics depends on the variability in the accumulation activity around the threshold. For the diffusive processes under consideration, this variability increases over time (as $\sim \sigma\sqrt{t}$). Therefore, assuming a constant noise amplitude implies that the CV increases as average reversal times lengthen, as illustrated in **Fig. 2.15**. If the noise becomes input-dependent, the shape of the distribution can be maintained even when reversal times change significantly (**Fig. 2.16**). Moreover, the correct dependence can be explicitly specified by the expression (2.22), so that drift and dispersion rates ν_{drift}^* and ν_{noise}^* change in the same proportions. For both modified WP and OU processes (see **Fig. 2.10** and **Fig. 2.14**), a scalar property is satisfied when operating in a DDR - i.e. distributions are shape-invariant. This constitute a necessary and sufficient condition for the scaling property to hold, and equivalently for the FPT densities to allow large

deviations of its mean, but not its shape, under variations of stimulus strength. In other words, by injecting more noise in faster accumulations, and less noise in slower accumulation, the balance between deterministic and stochastic factors can be maintained, while average dominance durations vary over the behavioural range observed in experiments.

Note that as previously mentioned (see end of Section 2.2.3), this can be done in the case of Gaussian diffusion models by considering the rescaling of the unit of time with stimulus strength x_{in} : for a Gaussian drift-diffusion process, the transformation $t \rightarrow x_{in} t$ would introduce a scaling factor $\sqrt{x_{in}}$ for the diffusion term.

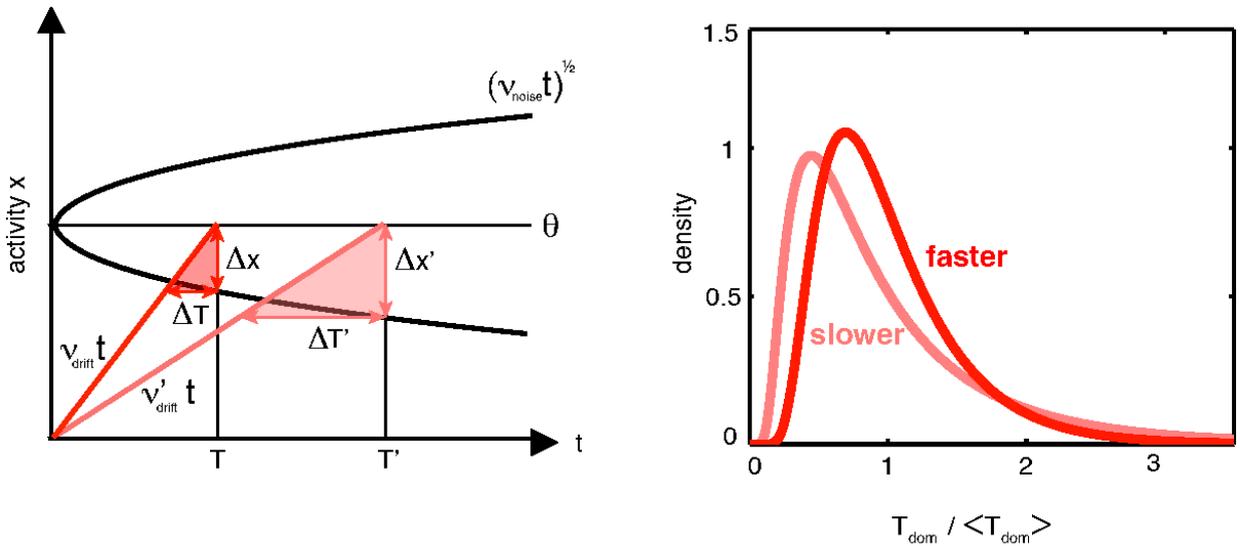


Figure 2.15: Schematic mechanism for the scaling property. Mean activity $x(t)$ accumulates approximately as $\nu_{drift} t$ (faster) or $\nu'_{drift} t$ (slower), and variance of activity as $\nu_{noise} t$ (faster) or $\nu'_{noise} t$ (slower), until threshold θ is reached. Coefficients of variation are approximately the same for time-to-threshold and for activity-to-threshold: $\Delta T/T \approx \Delta x/\theta$. (Left): For slower drift $\nu'_{drift} = \nu_{drift}/2$, but unchanged noise $\nu'_{noise} = \nu_{noise}$, activity-to-threshold is more variable, $\Delta x' = \Delta x\sqrt{2}$. When noise is reduced in the same proportion, $\nu'_{noise} = \nu_{noise}/2$, the variability of activity-to-threshold remains unchanged, $\Delta x' = \Delta x$. (Right): With constant noise, FPT distributions have different shapes for different input-levels.

Fluctuations are ubiquitous in neural systems, from the single-cell level up to the behavioural level. The common stance with regards to such fluctuations

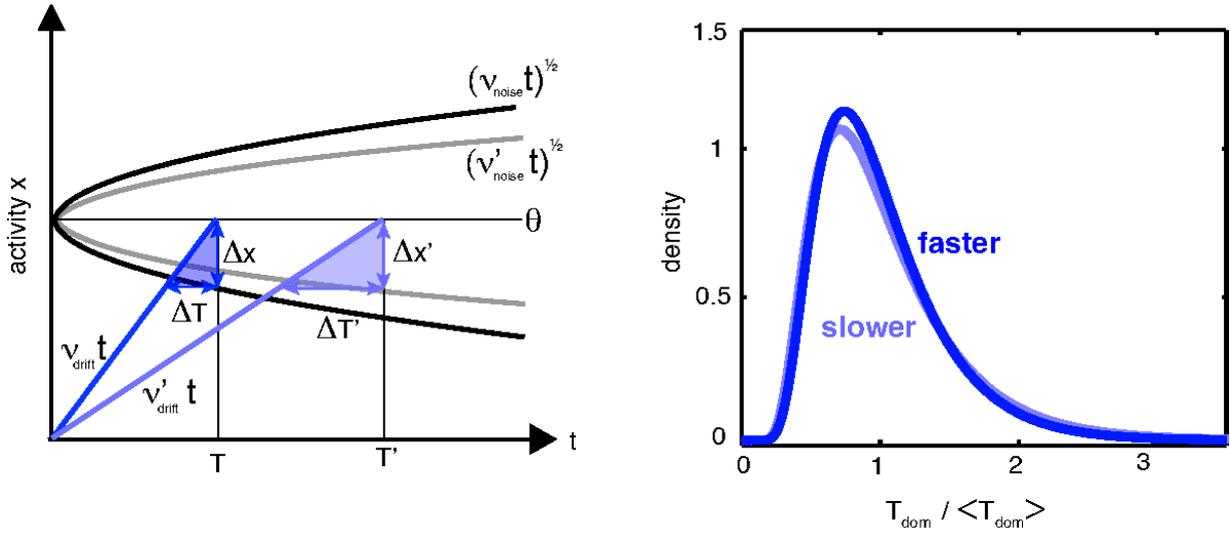


Figure 2.16: Schematic mechanism for the scaling property. Mean activity $x(t)$ accumulates approximately as $\nu_{drift}t$ (faster) or $\nu'_{drift}t$ (slower), and variance of activity as $\nu_{noise}t$ (faster) or $\nu'_{noise}t$ (slower), until threshold θ is reached. Coefficients of variation are approximately the same for time-to-threshold and for activity-to-threshold: $\Delta T/T \approx \Delta x/\theta$. (Left) For slower drift $\nu'_{drift} = \nu_{drift}/2$, but unchanged noise $\nu'_{noise} = \nu_{noise}$, activity-to-threshold is more variable, $\Delta x' = \Delta x\sqrt{2}$. (Left) When noise is reduced in the same proportion, $\nu'_{noise} = \nu_{noise}/2$, the variability of activity-to-threshold remains unchanged, $\Delta x' = \Delta x$. (Right): With input-dependent noise, FPT distributions have different shapes for different input-levels.

usually is to first disregard them, and to consider a deterministic model describing the average dynamics (mean-field or rate models). Noise is later introduced as an external parameter to account for the observed variability, and is tuned to optimally match empirical observations. Several studies have pointed out the unrealistic nature of a constant-noise assumption, in the inter-spike interval statistics of single-neuron models (see [Lánský and Sacerdote, 2001, Sacerdote and Lansky, 2002]), or in finite-size networks of spiking neurons, where fluctuations in activity depend on the average firing-rate (see [Mattia and Del Giudice, 2002, Buice and Chow, 2013, Deger et al., 2014]). A recent study investigating scale-invariant fluctuations in diverse multistable situations, such as binocular rivalry or in electroencephalogram recordings of the alpha-rhythm in humans, also highlighted similar plausible functional reasons why fluctuations should scale with input-strength or even with activity (e.g. state-dependent or multiplica-

tive noise) [Freyer et al., 2012]. Specifically, that allowing transitions between distinct states of activity whilst maintaining the CV of behavioural responses constant may reflect the need for our visual system to represent uncertainty adaptively.

In Chapter 2, we have shown that the scaling property could be obtained only if, during the stochastic accumulation of activity leading to perceptual reversals, modulations of input strength affect accumulation and dispersion rates proportionally. However, this dependence had to be artificially imposed. This motivates the introduction a more ‘physical’ hypothesis to explain the origin of this dependence.

Presumably, both deterministic and stochastic components of the accumulating process emerge as a collective property of fluctuating neural activity, which could help to constrain neural dynamics to a more detailed level. Furthermore, the scaling property is an extremely general and consistent observation, found to be valid across observers, displays types and stimulus qualities. A model which comprehensively accounts for this important empirical result may thus provide important insights into fundamental neural processes underlying perceptual decision.

This led to the following question: is there a class of stochastic processes which casually features such input-dependence of both the drift and the noise parameter? As mentioned in [Lánský and Sacerdote, 2001], a natural way to ensure that drift and dispersion rates remain proportional under varying stimulus condition would be to consider the *discrete* equivalent processes which Gaussian diffusion processes such as the WP and OU processes approximate. Specifically, we propose that these continuous random walks are in fact performed by the collective dynamics of a finite number of spontaneously active microscopic states. In this paradigm, drift and dispersion share common physical origins: they both emerge at the population

level as the macroscopic expression of fluctuating activity at the level of individual microscopic states.

In this chapter, we introduce an example of such random walk, known as the generalized Ehrenfest urn model. When operating far-from-equilibrium, first-passage-times of this process distribute exactly as observed in experiments, contrary to all other stochastic processes scrutinized in this study.

3.1 Introduction

The scaling property of multistable perception comports two important aspects. The first aspect is that the shape of reversal time densities is invariant under a wide array of conditions, indifferently of observers, the type of display used, or stimulus qualities. In particular, these different factors can greatly affect the average reversal rate, whilst keeping the distribution shape unchanged. The second aspect is that reversal time densities obey a particular shape, corresponding closely to a Gamma distribution (up to the third moment, with $c_v \approx 0.6$ and $\gamma_1 \approx 2c_v$).

We have already mentioned a simple and physically motivated hypothesis to ensure that the first aspect is verified. Namely, we consider a birth-death process, in which accumulation of activity is performed collectively, by an ensemble of bistable units undergoing spontaneous transitions between active and inactive states. The discreteness of such a representation automatically keeps normalized moments constant. This is because when considering discrete and spontaneous events, such as Poisson events, changes in events rate affect all orders of their individual switching statistics. In terms of collective activity, both accumulation and dispersion rates are then dependent on input-levels.

The origins of the second aspect of the scaling property, i.e. specific shape of reversal time densities is much less intuitive. Gaussian diffusion processes for instance typically produce inverse Gaussian (or Wald) FPT distribution, with a longer tail ($\gamma_1 \approx 3c_v$). Even the discreteness of some processes, such as the Balanced Poisson

process does not guarantee FPTs to be Gamma-distributed. These issues will be the focus of Chapter 4, where we will propose additional comparative analyses.

In this chapter, we introduce a discrete birth-death process, the Ehrenfest urn model, which satisfies both aspects of the scaling property. We present new analytical results, the moments at all order of its FPT density, and identify the different operating regimes relevant to our study [Cao et al., 2014]. Additionally, these results are generalized to a more neurophysiologically plausible case, where we consider interactions within assemblies of bistable units.

3.2 Ehrenfest urn model

Introduced in 1907 as the famous ‘dog-flea’ model of diffusion [Ehrenfest, 1907], this stochastic process provides a minimal model for the statistical dynamics of many microscopic and mesoscopic systems [Karlin and McGregor, 1965, van Kampen, 1981, Cao et al., 2014]. It was originally created in the context of thermodynamics, to propose a microscopic approach to the problem of heat exchange between two thermostats.

Its first-passage-time properties have been previously studied by several authors, who derived compact closed-form expressions for the mean FPT in special cases (i.e. with particular start- and end-points, see [Bingham, 1991, Palacios, 1993]), and in a more general case (for arbitrary start- and end-points [Dette, 1994]). In this chapter, we derive an expression to recurrently obtain moments of the FPT at any order from previous moments, and explicitly compute the first four moments in the general case [Cao et al., 2014].

3.2.1 Definition

Following a previous study from Gigante [Gigante et al., 2009], we consider a population of N bistable units, each of which switches spontaneously between two distinct states of activity. The transitions are set by independent Poisson events with rates $\nu_+ = \nu_+(s)$ (activation) and $\nu_- = \nu_-(s)$ (inactivation), depending on the level of sensory input s (see **Fig. 3.1**). All orders of the individual switching

statistics are input-dependent, not only the average switching rate. This renders both accumulation and dispersion rate dependent on current activation x and input level s , with important consequences for FPT moments. Additionally, we assume that the ratio ν_+/ν_- grows monotonically with s , to ensure that the asymptotic activity $x_{in}(s) = \frac{\nu_+}{\nu_+ + \nu_-}$ of the population to monotonically increase with the input. Individually, these locally bistable attractors model the dynamics of idealised cortical columns, or clusters of recurrently connected neurons. In networks of spiking-neurons, recurrent interactions can allow global self-sustaining states of activity to be spontaneously accessed [Amit, 1995]. In such networks, as postulated by our bistable units, transition rates between active and inactive states depend on the strength of the incoming synaptic input Amit and Brunel [1997], Wang [2002], Major and Tank [2004]. Collectively, they can be thought of as a multi-modular, or clustered neural network [Litwin-Kumar and Doiron, 2012]. At a later stage of this work, we will provide more plausible implementation of such neural structure, in terms of a network of spiking neurons (see 4.5).

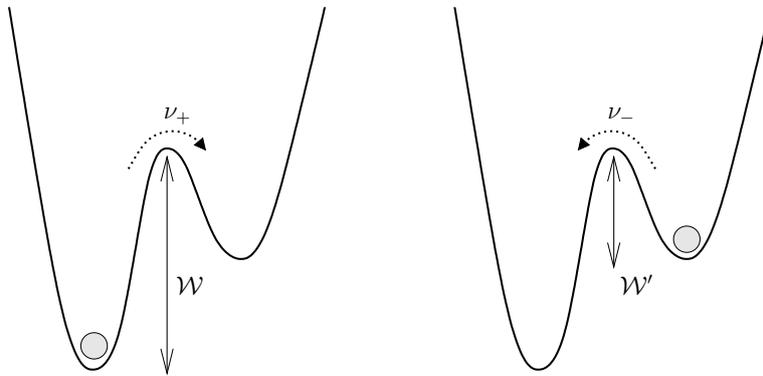


Figure 3.1: Schematic illustration of the individual switching properties between metastable states for a single bistable node in an *energy landscape* description. Upward transition rate ν_+ , and downward transition rate ν_- depend of the size of the energy barrier $\mathcal{W}(s)$, $\mathcal{W}(s)'$, which themselves depends on input strength s . This is analogous to Kramer's escape problem with two metastable states. Specifically, transition rates obey an Arrhenius-Van't Hoff relation where $\mathcal{W}(s)$, $\mathcal{W}(s)'$ are the corresponding activation energies.

To describe the collective dynamics of a pool of N independent nodes, one must

define the probability $P(n, t)$ of having n active nodes at a given time t . The variation $dP(n, t)$ over a time dt can be expressed as follows:

$$\begin{aligned} dP_n(t) = & + (N - n + 1)\nu_+ P_{n-1}(t)dt + (n + 1)\nu_- P_{n+1}(t)dt \\ & - (N - n)\nu_+ P_n(t)dt - n\nu_- P_n(t)dt \end{aligned} \quad (3.1)$$

Outward probability fluxes are counted negatively, and correspond to the transition $n \rightarrow n + 1$, where one of the $N - n$ inactive switches becomes active, or to the transition $n \rightarrow n - 1$, where one of the n active switches becomes inactive. Inward probability fluxes are counted positively, and correspond to the transition $n - 1 \rightarrow n$, where one of the $N - n + 1$ inactive switches becomes active, or to the transition $n + 1 \rightarrow n$, where one of the $n + 1$ active switches becomes inactive. This yields the following master equation where we identify upward transitions rate W_n^+ , and downward transition rates W_n^- :

$$\begin{aligned} \tau \frac{dP_n(t)}{dt} = & + \underbrace{[N - (n - 1)]\lambda}_{W_{n-1}^+} P_{n-1}(t) + \underbrace{(n + 1)\mu}_{W_{n+1}^-} P_{n+1}(t) \\ & - \left[\underbrace{(N - n)\lambda}_{W_n^+} + \underbrace{n\mu}_{W_n^-} \right] P_n(t) \end{aligned} \quad (3.2)$$

Here, we have defined $\lambda(s) = \frac{\nu_+}{\nu_+ + \nu_-} = x_{in}(s)$, the upward transition probability, $\mu(s) = 1 - \lambda = \frac{\nu_-}{\nu_+ + \nu_-}$, the downward transition probability, and $\tau(s) = \frac{1}{\nu_+ + \nu_-}$ the characteristic time-constant.

3.2.2 Detailed balance

The equilibrium distribution P_n^* for (3.2) can be obtained using the detailed balance relation:

$$W_n^- P_n^* = W_{n-1}^+ P_{n-1}^* \quad (3.3)$$

We can iterate (3.3) to obtain:

$$\begin{aligned}
 P_n^* &= P_0^* \prod_{k=1}^n \frac{W_{k-1}^+}{W_k^-} \\
 &= P_0^* \frac{(N - (n - 1)) \cdots N}{n \cdots 1} \lambda^n \mu^{-n} \\
 &= P_0^* \binom{N}{n} \lambda^n \mu^{-n}
 \end{aligned} \tag{3.4}$$

Using the fact that $\lambda + \mu = 1$, the normalization condition yields:

$$\begin{aligned}
 1 &= \sum_n P_n^* = \frac{P_0^*}{\mu^N} \sum_n \binom{N}{n} \lambda^n \mu^{N-n} \\
 P_0^* &= \mu^N
 \end{aligned} \tag{3.5}$$

We find that the equilibrium distribution is the binomial distribution with parameter λ .

$$P_n^* = \binom{N}{n} \lambda^n (1 - \lambda)^{N-n} \tag{3.6}$$

Equivalently, one can pick independent n Bernoulli processes with parameter λ (active nodes) and $N - n$ Bernoulli processes with parameter $1 - \lambda$ (the remaining inactive nodes), with the $\binom{N}{n}$ corresponding to the different ways to pick n active nodes in the population. Note that the detailed balance condition (3.3) is true for reversible Markov chains. This shows, a posteriori, that the GE process is reversible at equilibrium, since the correct stationary distribution has been obtained, as also shown in [Gigante et al., 2009]. Intuitively, this means that fluxes of upward and downward transitions are the same around equilibrium.

3.2.3 Mean and variance of the Ehrenfest process

1st-order: average dynamics

At first-order, equation (3.2) follows a simple exponential profile. We can derive the first-order solution for the average number of active nodes $\langle n \rangle$. Multiplying the

original master equation (3.2) by n and summing over n :

$$\begin{aligned} \tau \sum_n n \frac{dP(n,t)}{dt} &= \sum_n n [N - \underbrace{(n-1)}_k] \lambda P(n-1,t) + \sum_n n \underbrace{(n+1)}_k \mu P(n+1,t) \\ &\quad - \sum_n n [(N-n)\lambda + n\mu] P(n,t) \\ \tau \frac{d\langle n \rangle}{dt} &= \sum_k [N\lambda(k+1) - \lambda(k^2 - k) + \mu(k^2 - k) - N\lambda k + \lambda k^2 - \mu k^2] P(k,t) \\ &= N\lambda - \langle n \rangle \end{aligned}$$

Where the re-indexing $k \rightarrow k-1$ (in the first sum) and $k \rightarrow k+1$ (in the second sum) has been used to allow the factorization by $P(k,t)$. In terms of the the fraction of active switches in the pool ($x = n/N$), we obtain, for initial condition $x_0 = 0$:

$$\begin{aligned} \tau \frac{d\langle x \rangle}{dt} &= \lambda - \langle x \rangle \\ \langle x \rangle &= \lambda(1 - e^{-t/\tau}) \end{aligned} \quad (3.7)$$

2nd-order: variance

The same procedure can be followed to obtain the variance $\langle n^2 \rangle - \langle n \rangle^2$ of this process, this time, by calculating $\tau \sum_n n^2 \frac{dP(n,t)}{dt}$:

$$\tau \frac{d\langle n^2 \rangle}{dt} + 2\langle n^2 \rangle = 2N\lambda [N\lambda - \lambda + 1] - N\lambda [2N\lambda - 2\lambda + 1] e^{-\frac{t}{\tau}} \quad (3.8)$$

From the method of variation of the parameter, we look for a solution of the form:

$$\langle n^2 \rangle = C_0 e^{-\frac{2t}{\tau}} - N\lambda [2N\lambda - 2\lambda + 1] e^{-\frac{t}{\tau}} + N\lambda [N\lambda - \lambda + 1] \quad (3.9)$$

The integration constant C_0 can be obtained from the initial condition $\langle n^2 \rangle(0) = 0$:

$$C_0 = (N\lambda)^2 - N\lambda^2 \quad (3.10)$$

Finally, using (3.9) and (3.10), the solution is (see also [Pribram, 1994] p274 for the case $n(0) \neq 0$):

$$\begin{aligned}\langle n^2 \rangle - \langle n \rangle^2 &= N\lambda\mu + N\lambda(\lambda - \mu)e^{-\frac{t}{\tau}} - N\lambda^2e^{-\frac{2t}{\tau}} \\ \langle x^2 \rangle - \langle x \rangle^2 &= \frac{1}{N} \left[\lambda\mu + \lambda(\lambda - \mu)e^{-\frac{t}{\tau}} - \lambda^2e^{-\frac{2t}{\tau}} \right]\end{aligned}\quad (3.11)$$

From (3.7) and (3.11), we can compute the drift and dispersion rates:

$$\begin{aligned}\nu_{drift} &= \frac{\lambda}{\tau}e^{-\frac{t}{\tau}} \\ \nu_{noise} &= \frac{\lambda}{\tau N}e^{-\frac{t}{\tau}} \left[2\lambda e^{-\frac{t}{\tau}} - (\lambda - \mu) \right]\end{aligned}\quad (3.12)$$

In a DDR, we see that the GE process should satisfy the scaling property as the input strength λ varies:

$$\begin{aligned}\nu_{drift}^* &= \frac{\lambda}{\tau} \\ \nu_{noise}^* &= \frac{\lambda}{\tau N} \Rightarrow c_v^* = \frac{1}{\sqrt{\theta N}}\end{aligned}\quad (3.13)$$

3.2.4 General Solution

Equation (3.2) can be solved explicitly by introducing the generating function $G_i(z, t)$ of probability density $P_{ij}(t)$, the probability to be in j at time t , starting from i at $t = 0$:

$$G_i(z, t) = \sum_{k=0}^N P_{ik}(t)z^k \quad (3.14)$$

Importantly, the derivative of $G_i(z, t)$ with respect to z yields:

$$\frac{\partial}{\partial z} G_i(z, t) = \sum_{k=0}^N k P_{ik}(t)z^{k-1} \quad (3.15)$$

Therefore, defining the characteristic time $\tau = \frac{1}{\nu_+ + \nu_-}$ as well as upward transition probability $\lambda = \frac{\nu_+}{\nu_+ + \nu_-}$ and downward transition probability $\mu = \frac{\nu_-}{\nu_+ + \nu_-}$

$$\tau \frac{\partial}{\partial t} G_i(z, t) + (\lambda z + \mu)(z - 1) \frac{\partial}{\partial z} G_i(z, t) = N\lambda(z - 1)G_i(z, t) \quad (3.16)$$

This effectively converted the Master Equation (3.2) into a partial differential equation on the generating function $G_i(z, t)$ which we can solve using the method of characteristics (see for instance [Cox and Miller, 1972]). Characteristic equations are:

$$\frac{dt}{\tau} = \frac{dz}{(\lambda z + \mu)(z - 1)} = \frac{dG_i}{N\lambda(z - 1)G_i} \quad (3.17)$$

Integrating the first characteristic equation in (3.17) gives us an appropriate change of variable $\xi(z, t) = e^{-t/\tau} \frac{z-1}{\lambda z + \mu}$, so that an arbitrary function $f(\xi)$ satisfies the homogeneous equation associated with (3.16), i.e. with the right hand term set to zero. The general solution can be found, via the methods of variation of parameters, to be of the following form:

$$G_i(z, t) = (\lambda z + \mu)^N f_i \left[e^{-t/\tau} \frac{1 - z}{\lambda z + \mu} \right] \quad (3.18)$$

Where f_i is to be determined from the initial condition $G_i(z, 0) = z^i$. This corresponds to the initial state $P_i(0) = 1$, i.e. starting with the i -th node active, and all others, inactive:

$$f_i(x) = (1 + \lambda x)^{N-i} (1 - \mu x)^i \quad (3.19)$$

Assuming initial condition $i = 0$, i.e. starting without any active nodes, the generating function $G_0(z, t)$ writes:

$$\begin{aligned} G_0(z, t) &= (\mu + \lambda z)^N \left[1 - \frac{\lambda(z-1)}{\mu + \lambda z} e^{-t/\tau} \right]^{N-n_0} \\ &= [(\lambda e^{-t/\tau} + \mu) + \lambda(1 - e^{-t/\tau})z]^N \\ &= \sum_{k=0}^N \binom{N}{k} (\lambda e^{-t/\tau} + \mu)^{N-k} (\lambda(1 - e^{-t/\tau}))^k z^k \end{aligned} \quad (3.20)$$

Finally, we can identify the probability $P_{0k}(t) \doteq P_k(t)$ in (3.20):

$$P_k(t) = \binom{N}{k} (\lambda(1 - e^{-t/\tau}))^k (\lambda e^{-t/\tau} + \mu)^{N-k} \quad (3.21)$$

For $t \rightarrow \infty$, we retrieve the stationary distribution P_k^* (see (3.6)):

$$P_k^* = \binom{N}{k} \lambda^k (1 - \lambda)^{N-k} \quad (3.22)$$

The general solution $P_{ik}(t)$ for an initial condition different from 0 is given by (see detail of derivation in Appendix A):

$$P_{ik}(t) = \binom{N}{k} \frac{\lambda^k}{\mu} \sum_{n=0}^N \binom{N}{n} \lambda^n \mu^{N-n} K_i(n) K_k(n) e^{-\frac{nt}{\tau}} \quad (3.23)$$

Where K_n 's are Krawtchouk's orthogonal polynomials, defined as a special case of Gauss' confluent hypergeometric function ${}_2F_1$

$$K_n(x) = {}_2F_1(-n, -x, -N, 1/\lambda) = \sum_{k=0}^n \frac{(-n)_k (-x)_k}{(-N)_k k!} \left(\frac{1}{\lambda}\right)^k \quad (3.24)$$

$$(3.25)$$

Here, we have defined the Pochhammer symbol, or the rising factorial, as:

$$(a)_k = a(a+1) \dots (a+k-2)(a+k-1) \quad (3.26)$$

$$\begin{aligned}
K_0(x) &= 1 \\
-x\tau K_0(x) &= -\lambda_0 K_0(x) + \lambda_0 K_1(x) \\
-x\tau K_n(x) &= \mu_n K_{n-1}(x) - (\lambda_n + \mu_n) K_n(x) + \lambda_n K_{n+1}(x) \quad n \leq N \quad (3.29)
\end{aligned}$$

This three-terms recurrence is obeyed by the confluent hypergeometric function ${}_2F_1(-n, -x, -N; 1/\lambda)$, which can be verified by using one of Gauss' contiguous relations:

$$a({}_2F_1(a+1, b, c; z) - {}_2F_1(a, b, c; z)) = \frac{(c-a){}_2F_1(a-1, b, c; z) + (a-c+bz){}_2F_1(a, b, c; z)}{1-z} \quad (3.30)$$

Because in our case, N is finite, the hypergeometric function reduces to a Krawtchouk polynomial (see (3.25)). The Laplace transform \hat{F}_{ij} of the FPT distribution F_{ij} (starting in i and arriving for the first time in j , also assuming $i < j$) is [Karlin and McGregor, 1959]:

$$\hat{F}_{ij}(x) = \frac{K_i(-x\tau)}{K_j(-x\tau)} \quad (3.31)$$

Moments of the FPT distribution are obtained by calculating successive derivatives of (3.31), evaluated in zero:

$$\langle T_{ij}^n \rangle = (-1)^n \hat{F}_{ij}^{(n)}(0) \quad (3.32)$$

3.3.2 Moments formulae

From (3.32), the derivation of moments mostly relies on calculating the derivatives of the Pochhammer symbol (3.26), evaluated in zero. Up from the third order onwards, these expressions involve the generalised harmonic numbers $H_{n,r} = \sum_{k=1}^n k^{-r}$ (detailed derivations are presented in Appendix A).

$$(x)_k^{(1)} \xrightarrow{x \rightarrow 0} (k-1)! \quad (3.33)$$

$$(x)_k^{(2)} \xrightarrow{x \rightarrow 0} 2(k-1)!H_{k-1,0} \quad (3.34)$$

$$(x)_k^{(3)} \xrightarrow{x \rightarrow 0} 3(k-1)!(H_{k-1,0}^2 - H_{k-1,1}) \quad (3.35)$$

$$(x)_k^{(4)} \xrightarrow{x \rightarrow 0} 4(k-1)!(H_{k-1,0}^3 - 3H_{k-1,0}H_{k-1,1} + 2H_{k-1,2}) \quad (3.36)$$

After deriving these relations via direct calculations, we sought to obtain a general formula at any order. The generalisation is in fact straightforward. Indeed, the generating function for the Pochhammer symbol is known in terms of the Stirling numbers of the first kind, a combinatorial number, noted $\left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right]$ (not to be confused with the binomial coefficients):

$$(x)_k = \sum_{r=0}^k \left[\begin{smallmatrix} k \\ r \end{smallmatrix} \right] x^r \quad (3.37)$$

Taking the m -th derivative of (3.37), and taking the limit $x \rightarrow 0$ leaves us only the term of order m , so that:

$$\frac{d^m}{dx^m} (x)_k = m! \left[\begin{smallmatrix} k \\ m \end{smallmatrix} \right] \quad (3.38)$$

In particular, the Stirling number of the first kind in terms of the generalized harmonic numbers $H_{n,r}$:

$$\left[\begin{smallmatrix} k \\ m \end{smallmatrix} \right] = \frac{(k-1)!}{(m-1)!} w(k, m-1) \quad (3.39)$$

Where w 's are defined recursively by:

$$\begin{cases} w(k, 0) = 1 \\ w(k, m) = \sum_{k=0}^{m-1} (1-m)_k w(n, m-1-k) \end{cases} \quad (3.40)$$

Expression (3.40) indeed retrieves previously calculated expressions, (3.33), (3.34), (3.35), and (3.36). Also, formula (3.32) can be further expanded by using the n -th derivative of a quotient of function (as shown in [Xenophontos, 2007]). In our case, since $K_n(0) = 1$ for all n , all quotients have vanished from the final equation. The result is a recurrence formula for the n -th moment of the FPT distribution, given all prior moments:

$$\langle T_{ij}^n \rangle = (-\tau)^n \left(K_i^{(n)}(x) - K_j^{(n)}(x) \right) \Big|_{x=0} - \sum_{k=1}^{n-1} (-\tau)^k \binom{n}{k} K_j^{n-k} \langle T_{ij}^k \rangle \quad (3.41)$$

For the four first moments, we obtained:

$$\langle T_{ij} \rangle = \left(K_j^{(1)} - K_i^{(1)} \right) \Big|_{x=0} \quad (3.42)$$

$$\langle T_{ij}^2 \rangle = \left((K_i^{(2)} - K_j^{(2)}) + 2K_j^{(1)} \langle T_{ij} \rangle \right) \Big|_{x=0} \quad (3.43)$$

$$\langle T_{ij}^3 \rangle = \left((K_j^{(3)} - K_i^{(3)}) - 3K_j^{(2)} \langle T_{ij} \rangle + 3K_j^{(1)} \langle T_{ij}^2 \rangle \right) \Big|_{x=0} \quad (3.44)$$

$$\langle T_{ij}^4 \rangle = \left((K_i^{(4)} - K_j^{(4)}) + 4K_j^{(3)} \langle T_{ij} \rangle - 6K_j^{(2)} \langle T_{ij}^2 \rangle + 4K_j^{(1)} \langle T_{ij}^3 \rangle \right) \Big|_{x=0} \quad (3.45)$$

Using (3.41)-(3.45), we obtain the four first moments of the FPT distribution.

$$\langle T_{ij} \rangle = \tau \sum_{k=1}^j \frac{(-j)_k - (-i)_k}{k(-N)_k} \frac{1}{\lambda^k} \quad (3.46)$$

$$\langle T_{ij}^2 \rangle = \tau^2 \sum_{k=1}^j \frac{(-i)_k - (-j)_k}{k(-N)_k} \frac{2H_{k-1}}{\lambda^k} + 2\langle T_{ij} \rangle \langle T_{0j} \rangle \quad (3.47)$$

$$\begin{aligned} \langle T_{ij}^3 \rangle &= \tau^3 \sum_{k=1}^j \frac{(-j)_k - (-i)_k}{k(-N)_k} \frac{3(H_{k-1,1}^2 - H_{k-1,2})}{\lambda^k} \\ &\quad + 3(\langle T_{0j}^2 \rangle - 2\langle T_{0j} \rangle^2) \langle T_{ij} \rangle + 3\langle T_{0j} \rangle \langle T_{ij}^2 \rangle \end{aligned} \quad (3.48)$$

$$\begin{aligned} \langle T_{ij}^4 \rangle &= \tau^4 \sum_{k=1}^{\theta} \frac{(-i)_k - (-j)_k}{k(-N)_k} \frac{4(H_{k-1,0}^3 - 3H_{k-1,0}H_{k-1,1} + 2H_{k-1,2})}{\lambda^k} \\ &\quad + 4(\langle T_{0j}^3 \rangle - 6\langle T_{0j}^2 \rangle \langle T_{0j} \rangle + 6\langle T_{0j} \rangle^3) \langle T_{ij} \rangle \\ &\quad + 6(\langle T_{0j}^2 \rangle - 2\langle T_{0j} \rangle^2) \langle T_{ij}^2 \rangle + 4\langle T_{0j} \rangle \langle T_{ij}^3 \rangle \end{aligned} \quad (3.49)$$

3.3.3 Simulations

We can now express the mean μ_1 , coefficient of variation c_v , skewness γ_1 and kurtosis κ of the FPT density, as a function of its raw moments calculated in the previous section:

$$\mu_1 = \langle T_{ij} \rangle \quad (3.50)$$

$$c_v = \sqrt{\frac{\langle T_{ij}^2 \rangle}{\langle T_{ij} \rangle^2} - 1} \quad (3.51)$$

$$\gamma_1 = \frac{1}{c_v^3} \left[\frac{\langle T_{ij}^3 \rangle}{\langle T_{ij} \rangle^3} - 3c_v^2 - 1 \right] \quad (3.52)$$

$$\kappa = \frac{1}{c_v^4} \left[\frac{\langle T_{ij}^4 \rangle}{\langle T_{ij} \rangle^4} - 4\gamma_1 c_v^3 - 6c_v^2 - 1 \right] \quad (3.53)$$

Even higher-order moments can be worked out with comparable effort. The recursive expression (3.41), as well as higher-order moments (3.47),(3.48) and (3.49) are new. Only expression (3.46) for the first moment was published previously [Dette, 1994]. In **Fig. 3.2**, **Fig. 3.3** and **Fig. 3.4**, direct simulations of N stochastic bistable columns are compared to the predictions of calculated expressions (3.47),(3.48) and (3.49). Specifically, the mean μ_1 , coefficient of variation c_v , skewness γ_1 , and kurtosis κ of the FPT density are shown for different initial conditions $x_0 \in \{0, 0.2, 0.4\}$ and different values of ν_{\pm} . The latter values were chosen to set the steady-state activation x_{in} at various locations below or above the threshold θ/N ($x_{in} - \theta/N \in [-0.2, +0.2]$). Our analytical formulae are indeed exact and show a remarkable agreement with direct simulations.

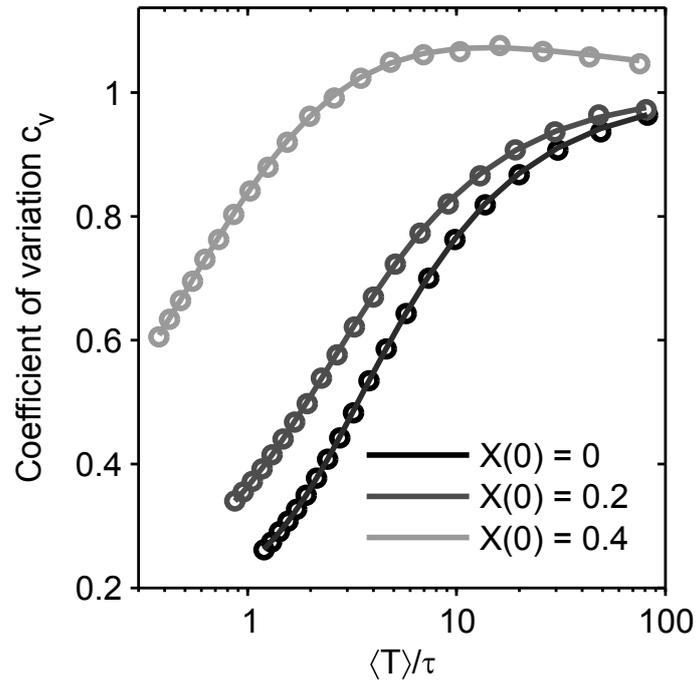


Figure 3.2: Agreement between the analytical expression of the coefficient of variation (3.51) of the FPT and simulated realisations of the process, as a function of the distribution's mean (3.50). Each symbol (circles) represents 160000 simulated FPTs, and transition rates ν_+ and ν_- were varied to span $\lambda - \theta/N \in [-0.2, +0.2]$. Threshold $\theta = 25$, number of nodes or columns $N = 50$, relaxation time $\tau = (\nu_+ + \nu_-)^{-1}$. *Reprinted from [Cao et al., 2014] with the co-authors' permission.*

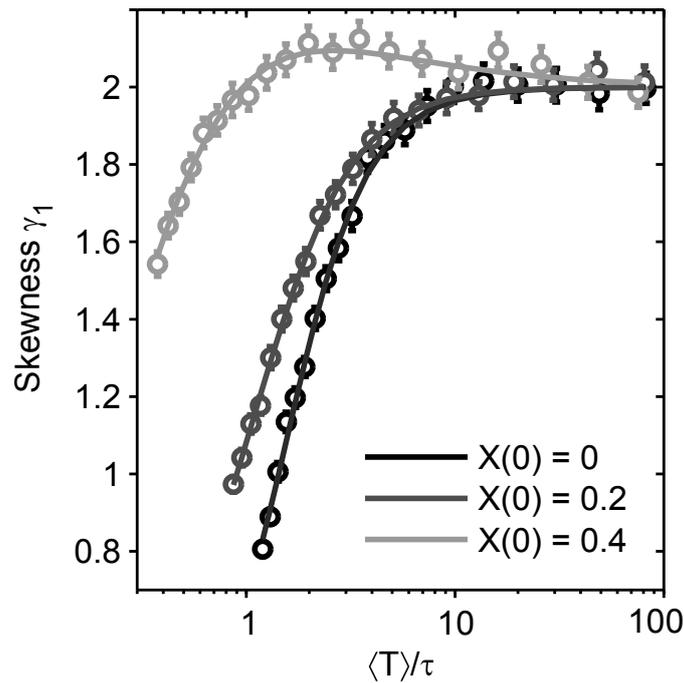


Figure 3.3: Agreement between the analytical expression of the skewness γ_1 (3.52) of the FPT and simulated realisations of the process, as a function of the distribution's mean (3.50). Each symbol (circles) represents 160000 simulated FPTs, with error bars indicating SEM. Transition rates ν_+ and ν_- were varied to span $\lambda - \theta/N \in [-0.2, +0.2]$. Threshold $\theta = 25$, number of nodes or columns $N = 50$, relaxation time $\tau = (\nu_+ + \nu_-)^{-1}$. Reprinted from [Cao et al., 2014] with the co-authors' permission.

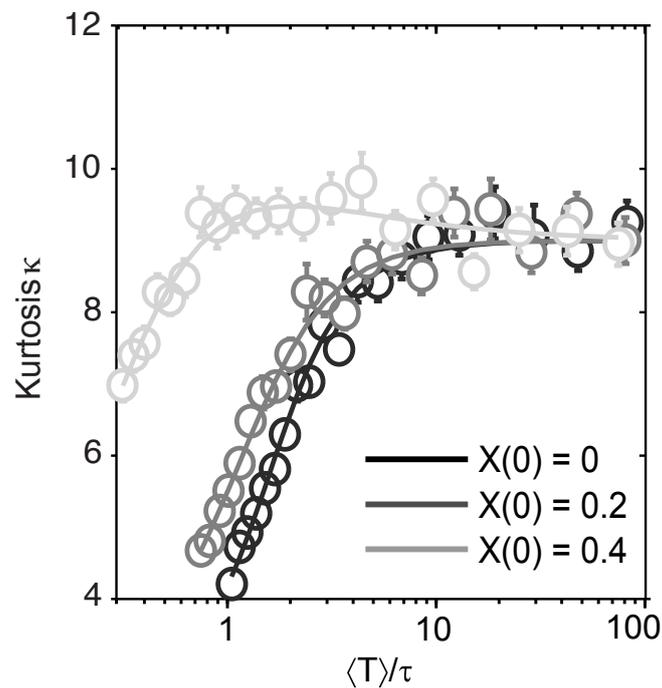


Figure 3.4: Agreement between the analytical expression of the kurtosis κ (3.53) of the FPT and simulated realisations of the process, as a function of the distribution's mean (3.50). Each symbol (circles) represents 160000 simulated FPTs, with error bars indicating SEM. Transition rates ν_+ and ν_- were varied to span $\lambda - \theta/N \in [-0.2, +0.2]$. Threshold $\theta = 25$, number of nodes or columns $N = 50$, relaxation time $\tau = (\nu_+ + \nu_-)^{-1}$.

3.3.4 Low-Threshold Approximation

In the low-threshold limit, so that $\theta \ll N$, we can get simple approximates for $\mathbb{E}[T_{0\theta}]$ and $\mathbb{E}[T_{0\theta}^2]$, and therefore, for the CV.

For the first moment:

$$\begin{aligned}
\mathbb{E}[T_{0\theta}] &= \sum_{k=1}^{\theta} \frac{(-\theta)_k}{(-N)_k} \frac{1}{k\nu_+^k \tau^{k-1}} \\
&= \frac{\theta}{N} \frac{1}{\nu_+} + \frac{\theta(\theta-1)}{N(N-1)} \frac{1}{2\nu_+^2 \tau} + \frac{\theta(\theta-1)(\theta-2)}{N(N-1)(N-2)} \frac{1}{3\nu_+^3 \tau^2} + \dots \\
&= \frac{\theta}{N\nu_+} \left(1 + \underbrace{\frac{\theta-1}{N(1-1/N)}}_{\sim N} \frac{1}{2\nu_+ \tau} + \underbrace{\frac{(\theta-1)(\theta-2)}{N^2(1-1/N)(1-2/N)}}_{\sim N^2} \frac{1}{3\nu_+^2 \tau^2} + \dots \right) \\
&\simeq \frac{\theta}{N\nu_+} \left(1 + \frac{\theta-1}{N} \frac{1}{2\nu_+ \tau} + \frac{(\theta-1)(\theta-2)}{N^2} \frac{1}{3\nu_+^2 \tau^2} + \dots \right) \tag{3.54}
\end{aligned}$$

Similarly, for the second moment and CV, using $\mathbb{E}[T_{0\theta}] = \langle T_{0\theta} \rangle = \frac{\theta}{N\nu_+}$:

$$\mathbb{E}[T_{0\theta}^2] \simeq \frac{\theta(\theta-1)}{(N\nu_+)^2} + \frac{(\theta-1)(\theta-2)}{(N\nu_+)^3 \tau} + \dots \tag{3.55}$$

$$CV[T_{0\theta}] \simeq \frac{1}{\sqrt{\theta}} \left(1 + \frac{\theta-1}{2\theta} \frac{\langle T_{0\theta} \rangle}{\tau} \right) \tag{3.56}$$

Once again, in the DDR, $\langle T_{0\theta} \rangle \ll \tau$, and we retrieve the previous approximated expression (3.13) by substituting $\theta \rightarrow \theta N$. When combining low-threshold and drift-dominated regime, the CV of FPTs depend exclusively on the value of the threshold, and increases as the mean FPT increases, when approaching the NDR.

3.4 Regime Identification

3.4.1 Drift- and noise-dominated regimens and the scaling property

Recall that steady-state activity $x_{in}(s)$ (indirectly) represents sensory input. For different combinations of x_{in} and threshold θ/N , the analysis reveals curves of

constant c_v with two branches (**Fig. 3.5**, left). This qualitative picture changes with neither population size N nor initial conditions (**Fig. 3.5**, right). To each of these branches correspond a distinct operating regime, which differ greatly, not only qualitatively, but also quantitatively.

To the lower, horizontal branch, corresponds a relatively low-threshold regime ('suprathreshold', or DDR $\theta < x_{in}$). In this regime, c_v remains nearly constant when the threshold θ/N is fixed slightly above baseline activity x_0 (red dashed line, **Fig. 3.6**, left). Importantly, the threshold can be set to a fixed value without violating the constrain $c_v \approx 0.6$ (red shadings, **Fig. 3.7**, left) To the higher, diagonal branch, correspond a high-threshold regime ('subthreshold', or NDR $\theta > x_{in}$). This reflects the situation where the threshold θ/N is set slightly above steady-state activity x_{in} (blue dashed line, **Fig. 3.6**, left). In this case, the threshold can not be set to a fixed value without violating the scaling property. Rather, it must be adjusted, for different values of x_{in} . Indeed, the distance-to-threshold (from the steady-state value) $\theta/N - x_{in}$ and must remain constant to satisfy $c_v \approx 0.6$ (blue shadings, **Fig. 3.7**, left).

3.4.2 Drift- and noise-dominated regimens and the distribution shape

Quantitatively, the two regimes differ greatly on several accounts: fluctuations far-from-equilibrium (i.e. in the DDR), yields significantly different FPT statistics than fluctuations around equilibrium (i.e. in the NDR). First, mean perceptual durations $\langle T \rangle$ are comparatively fast ($\langle T \rangle \ll \tau$) in the DDR regime – as only the stochastic accumulation of few additional active columns is needed in order to breach threshold – and comparatively slow ($\langle T \rangle \gg \tau$) in the NDR regime – as fluctuations must carry activity above the steady-state (**Fig. 3.7**, left). A telling difference concerns the distribution shape (**Fig. 3.6**, right), with FPT distribution for the DDR being significantly less skewed than in the NDR, for the same value of $c_v \in [0.5, 0.6]$.

The difference between both distribution can be made particularly salient by com-

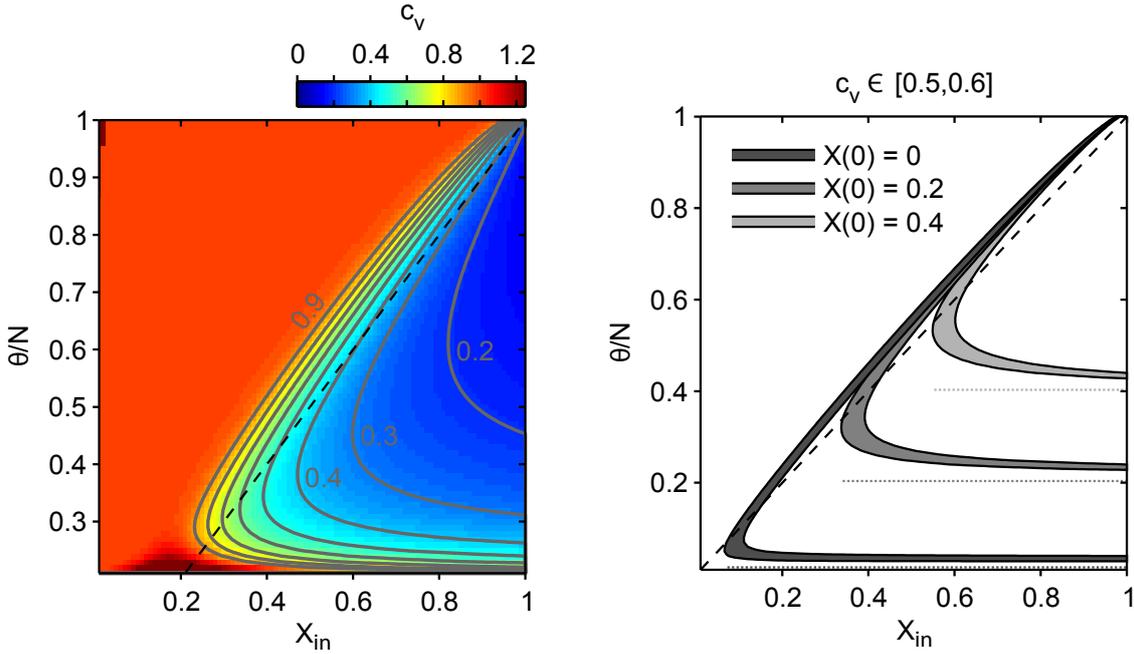


Figure 3.5: Areas of with near-constant coefficient of variation $c_v \approx 0.6$. (*Left*): Coefficient of variation c_v of FPT distribution as a function of the steady state $x_{in} = \lambda$ (input strength) and perceptual threshold θ . Curves with identical c_v are marked in gray. Population size $N = 100$ and initial condition $x_0 = 0.2$. (*Right*): Ranges with $c_v \in [0.5, 0.6]$ for different initial conditions $x_0 = 0, 0.2, 0.4$ (from dark to light shades, respectively). Dotted lines mark values of x_0 , dashed line marks steady-state activity Nx_{in} . *Reprinted from [Cao et al., 2014] with the co-authors' permission.*

paring the range of values for their skewness (**Fig. 3.7**, right). Whereas the FPT distribution in the DDR corresponds to a Gamma distribution up to the third moment ($\gamma_1 \simeq 2c_v$), the NDR regime exhibits a refractory period of $\simeq \tau$ and a heavier, exponential-like, tail (with $\gamma_1 \gtrsim 3c_v$), characteristic of an inverse Gaussian distribution. The reason is that the threshold is reached in two separate phases: first, activity must reach the steady-state level and, second, so that fluctuations around the steady-state can allow the system to cross the threshold. The second step is similar to the - solely noise driven - escape from a potential-well [Kramers, 1940, Hanggi et al., 1990].

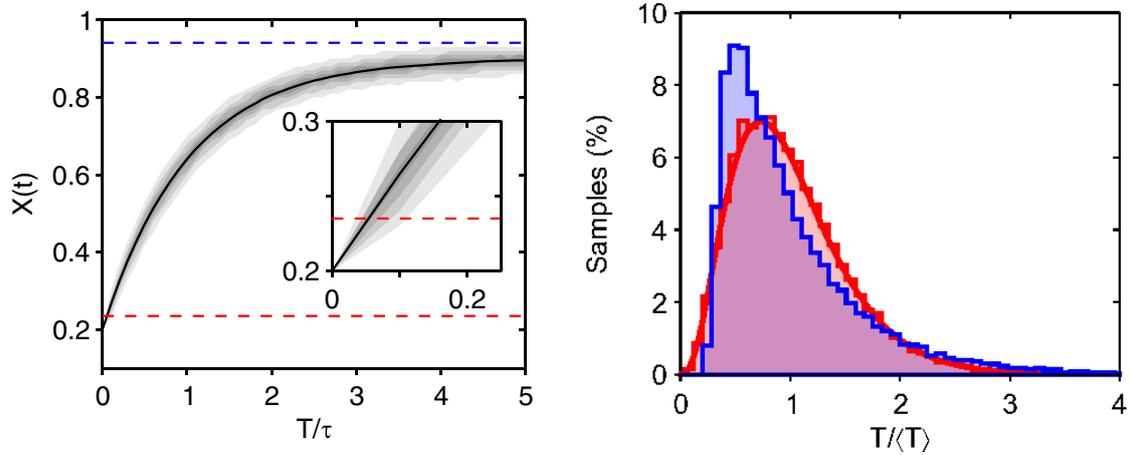


Figure 3.6: Operating regime for the scaling property. (*Left*): Time evolution of density $P_n(t)$ for $x_0 = 0.2$ and x_{in} [20 000 simulated realizations of $x(t)$]. Shadings indicate 10 percentile steps. The inset magnifies the initial dynamic. Either a suprathreshold (dashed red) or a subthreshold (dashed blue) yields a $c_v = 0.55$. (*Right*): The FPT distribution is significantly less skewed in the suprathreshold (red histogram) than in the subthreshold regime (blue histogram). Reprinted from [Cao et al., 2014] with the co-authors' permission.

3.4.3 Interacting pool

Linear self-excitation

Throughout this chapter, we have considered the collective accumulation of activity in a finite assembly of bistable units. Additionally, we have assumed such units were *independent*, for the sake of simplicity, so that exact analytical results could be derived. In terms of neural populations, these units can be thought of as idealized cortical columns, supporting a given perceptual appearance, or representation, of a stimulus. From this perspective, assuming independence between units of the same assembly would imply that they would correspond to unstructured representations of independent stimulus features. To demonstrate the generality of our results, we now relax this questionable assumption and consider an assembly of interacting units. Recurrent interactions within the assembly of bistable units will provide additional input to individual element, making transition rates ν_{\pm} dependent on

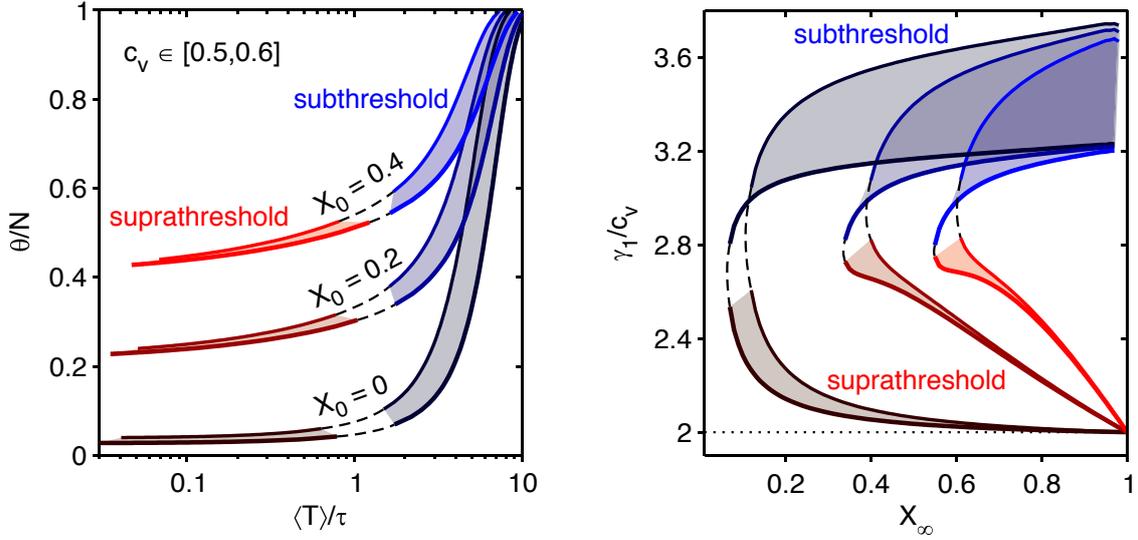


Figure 3.7: (Left): Threshold values θ/N required to obtain $c_v \in [0.5, 0.6]$ with different mean FPTs $\langle T \rangle$ in the supra- and subthreshold regimes (red and blue, respectively). Shading represents different values of $x_0 \in \{0, 0.2, 0.4\}$. (Right): Ratio between skewness γ_1 and c_v of FPT distribution for ranges with $c_v \in [0.5, 0.6]$ and different values of $x_0 \in \{0, 0.2, 0.4\}$. Reprinted from [Cao et al., 2014] with the co-authors' permission.

collective activity $x(t)$ in a mean-field approximation.

In a strong DDR, where x_{in} is high, and θ is low, it is sufficient to consider a linear dependence of excitatory couplings:

$$\nu_+(x) = \nu_+^{(0)} + \nu_+^{(1)}x(t) \quad (3.57)$$

Since in this regime, $x(t)$ remains mostly in a narrow range just below threshold (see for example inset in **Fig. 3.6**, left), the following results would pertain even if higher-order terms of the Taylor expansion were to be included in Eq. (3.57). This is because greater modulations of $x(t)$ would be necessary for non-linear effects to be noticeable, which is not permitted in such a restricted range between initial condition and threshold. Interestingly, greater coupling strength $\nu_+^{(1)}$ leaves the c_v of the FPT essentially unchanged (see **Fig. 3.8 A**, top). Coupling decreases mean FPTs, but mostly in the DDR regime (see **Fig. 3.8 A**, bottom), suggesting that such coupling affects the dynamics far-from-equilibrium more than it affects the dynamics around equilibrium.

Such coupling introduces modulations of both the asymptotic value x_{in} and the characteristic time-scale τ . In the NDR, a sufficient increase of x_{in} could effectively push the system across the bifurcation to the DDR, so we compensate the changes in x_{in} due to increasing coupling strength by increasing the downward switching rate $\nu_-^{(0)}$. In the noise-dominated case, deviations of the mean FPT are less important, because significant changes in the time-scale of the system have already occurred during the approach to equilibrium, while in the drift-dominated case, changes of time-scales occur as the system approaches and crosses the threshold.

Remarkably, the difference between the distribution shapes between DDR and NDR, as expressed in the ratio γ_1/c_v is invariant with coupling strength and initial conditions (see **Fig. 3.8 B**). Recurrent interactions between bistable units shift the FPT distribution in time, *without changing its shape*. In effect, interactions scale transitions rates ν_{pm} just as a progressive compression of the unit of time would do.

Time-scale modulation

To understand this compression of time-scales, we consider the effect of the coupling (3.57) on the deterministic component of the dynamics. For $\nu_{\pm}^{(1)} = 0$, we have a simple exponential relaxation to asymptotic value $\lambda_0 = \nu_+^{(0)}/(\nu_+^{(0)} + \nu_-^{(0)})$, with characteristic time-scale $\tau_0 = (\nu_+^{(0)} + \nu_-^{(0)})^{-1}$, so that the solution for the initial condition $x(t=0) = 0$ is:

$$x(t) = \lambda_0(1 - e^{-t/\tau_0}) \quad (3.58)$$

Therefore, we introduce the linear coupling into , which becomes a Riccati equation with constant coefficients depending on the non-interacting transition rates $\nu_{\pm}^{(0)}$ and the recurrent coupling rates $\nu_{\pm}^{(1)}$:

$$\dot{\tilde{x}} = \nu_+^{(0)} - (\nu_+^{(0)} + \nu_-^{(0)} - \nu_+^{(1)})\tilde{x} - (\nu_+^{(1)} + \nu_-^{(1)})\tilde{x}^2 \quad (3.59)$$

We introduce x_{\pm} , the conjugate roots of the second-order polynomial $Q(x) = ax^2 + bx + c = a(x - x_+)(x - x_-)$, where: $a = -(\nu_+^{(1)} + \nu_-^{(1)})$, $b = -(\nu_-^{(0)} + \nu_-^{(0)} - \nu_+^{(1)})$,

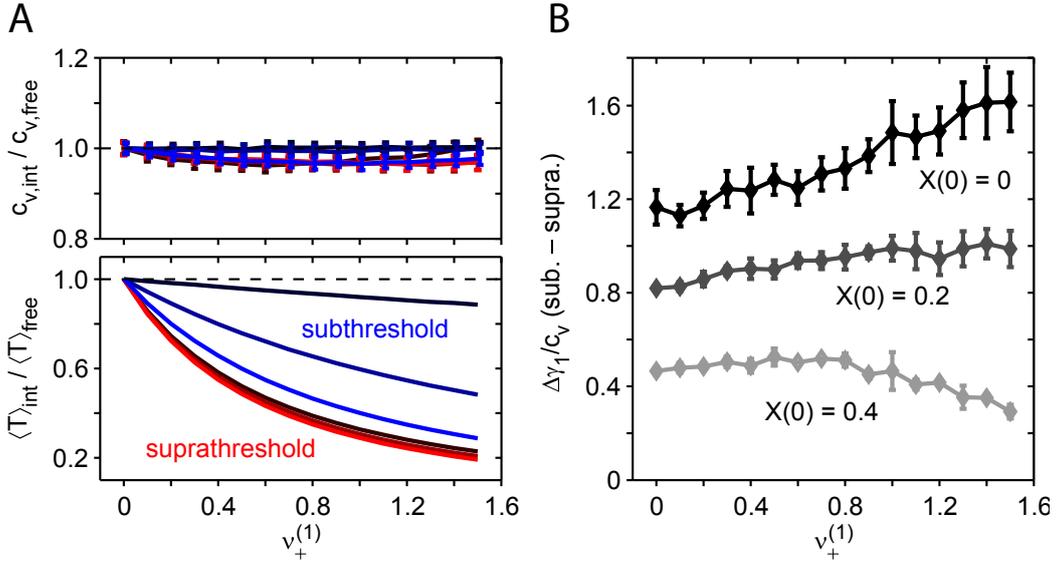


Figure 3.8: Higher-order statistics of FPT for populations of interacting bistable units. **A (Top):** comparison of coefficients of variation for interacting units, $c_{v,int}$, and independent units, $c_{v,free}$, for various coupling values $\nu_+^{(1)}$. **(Bottom):** comparison of mean FPT for the interacting units, $\langle T_{int} \rangle$ and the independent units $\langle T_{free} \rangle$. Colors distinguish supra- and subthreshold regimes and hues indicate different initial conditions x_0 . **B** Difference in the shape of FPT distribution shape between sub- and suprathreshold regimes, as expressed by $\frac{\Delta\gamma_1}{c_v}$. Shading distinguishes initial conditions. The baseline rates $\nu_{\pm}^{(0)}$ were chosen to obtain $x_{in} = 0.8$ and $\tau = 1$ with zero coupling (i.e. $\nu_+^{(1)} = 0$). In the subthreshold regime, $\nu_+^{(1)} > 0$ were compensated by increasing the baseline inactivation $\nu_-^{(0)}$, such as to maintain $x_{in} = 0.8$. For each dynamical regime, θ was set to maintain $c_{v,free} \approx 0.6$. Symbols represents mean and SEM from 5000 simulated realizations. *Reprinted from [Cao et al., 2014] with the co-authors' permission.*

$c = \nu_+^{(0)}$, as well as the discriminant $\Delta = \sqrt{b^2 - 4ac}$. Using the factorised form of the polynomial Q , we can integrate (3.59) by separation of variables with initial condition $x(t=0) = 0$:

$$\tilde{x}(t) = x_- \frac{1 - e^{-\sqrt{\Delta}t}}{1 - \frac{x_-}{x_+} e^{-\sqrt{\Delta}t}} \quad (3.60)$$

We can easily verify that $(\sqrt{\Delta})^{-1} \xrightarrow{\nu_{\pm}^{(1)} \rightarrow 0} \tau_0$, so that $(\sqrt{\Delta})^{-1}$ is the new characteristic time-scale for the coupled system. Similarly, the root x_- corresponds to the

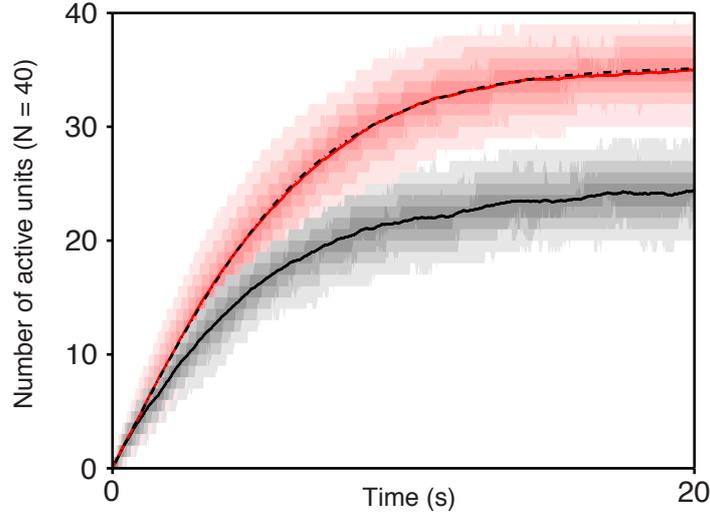


Figure 3.9: Illustration of the dynamics for a linearly coupled system $\nu_{\pm} = \nu_{\pm}^{(0)} + \nu_{\pm}^{(1)} x$, with $\nu_{+}^{(1)} = 0.2$ and $\nu_{+}^{(0)} = 0$. Shaded traces correspond to the evolution of the probability density for the fluctuating system (red: linearly coupled, black: uncoupled). Thick lines show the average of the probability densities at each time-step (red: linearly coupled, black: uncoupled). Dotted lines correspond to the solution obtained in (3.60)

new asymptotic value for the coupled system:

$$\begin{aligned} x_- &= -\frac{\nu_{+}^{(0)}}{\nu_{+}^{(1)} + \nu_{-}^{(1)}} \frac{1}{x_+} \\ &= 2 \frac{\nu_{+}^{(0)}}{\nu_{+}^{(0)} + \nu_{-}^{(0)} - \nu_{+}^{(1)} + \sqrt{\Delta}} \xrightarrow{\nu_{\pm}^{(1)} \rightarrow 0} \lambda_0 \end{aligned}$$

The dynamics of the coupled system can be envisaged in two different ways. A first interpretation is to consider an exponential relaxation toward a moving asymptote value $x_-(t) \doteq x_- \left(1 - \frac{x_-}{x_+} e^{-\sqrt{\Delta}t}\right)^{-1}$, for a fixed time-scale $\sqrt{\Delta}^{-1}$ as in (3.61). A second interpretation, which is more relevant for us, is to consider an exponential relaxation toward a fixed asymptote x_- , for a time-modulated time-scale $\tau(t)$, as in (3.62).

$$\tilde{x}(t) = x_-(t)(1 - e^{-t\sqrt{\Delta}}) \quad (3.61)$$

$$= x_-(1 - e^{-t/\tau(t)}) \quad (3.62)$$

We can identify the function $\tau(t)$ by inverting equation (3.62). We find a strictly decreasing function of t , converging to the time-scale of the coupled system $(\sqrt{\Delta})^{-1}$. As the activity rises, the characteristic time of the system progressively diminishes:

$$\tau(t) = \frac{1}{\sqrt{\Delta}} \left[1 + \frac{1}{t\sqrt{\Delta}} \ln \left(\frac{x_+ - x_- e^{-t/\tau}}{x_+ - x_-} \right) \right]^{-1} \quad (3.63)$$

To summarize, introducing excitatory couplings in assemblies of bistable units, only affects the mean FPT, and the shape of the FPT distribution is left essentially invariant, in both drift- and noise-dominated regimens. In particular, stronger interactions reduce the characteristic time-constant of the assembly, and yield stronger modulations of the mean FPT in when the system is far-from-equilibrium (DDR) than around equilibrium (NDR).

3.5 Summary

In this chapter, we have introduced a stochastic process belonging to different class of random walks, i.e. a birth-death process, also known as the generalized Ehrenfest process. We established new closed-form formulae for the first four moments of its FPT distribution [Cao et al., 2016]. Note that only the first order (average FPT) had previously been published [Dette, 1994].

We also performed a regime identification which revealed two separate regimens for which the scaling property is obtained, i.e. CV remains constant at $c_v \approx 0.6$ while the distribution mean can shifts by an order of magnitude with sensory input. These two regimens correspond to two threshold settings, corresponding to two different dynamics in approaching and crossing the threshold (DDR and NDR). Importantly, only in the DDR, is the characteristic Gamma-like shape reported in a wide range of behavioural experiments also obtained.

A low and fixed threshold, placed just above baseline activation ($\theta/N - X_0 \ll 1$), produces short mean durations (fractions of the relaxation time τ) and a Gamma-like skewness $\gamma_1 \approx 2c_v$, precisely matching experimental observations (see **Fig. 4.2**). A high and variable threshold, placed just above the steady-state

activation ($\theta/N - X_\infty \ll 1$), predicts far longer mean durations (multiples of τ) and significantly higher skewness $\gamma_1 \geq 3c_v$, at odds with experimental evidence.

The scaling property and the shape of the FPT distribution therefore depends on the collective properties of the population of bistable units, as well as the population size N and threshold placement θ . In contrast, the mean FPT reflects the intrinsic dynamics of individual columns, and in particular, i.e. the relaxation time τ . This additional degree of freedom may also help to explain the spurious variation of reversal times between observers, as different observers typically have different mean switching rates.

We conclude that such discrete accumulation of activity using a finite population of bistable units, operating far-from-equilibrium, naturally explains both the scaling property of multistable perception and the characteristic Gamma shape of reversal time densities observed experimentally.

Scaling Property and Distribution Shape

In Chapter 2, we have established one of the key requirements for a scaling property to be obtained. Namely, that the deterministic and stochastic components of the random walk underlying perceptual reversals must obey and maintain a specific balance. Therefore, a certain *input-dependence* of the noise appeared necessary and was explicitly derived.

In Chapter 3, we have shown that the GE process could reproduce both the scaling property *and* the distribution shape. This process accumulates activity with bistable units, which ensures the scaling property. This is because the individual dynamics of each bistable unit corresponds to the statistics of Poisson events, whose rate depends on the input-levels. When input-levels change, not only does the average time between events changes, but also its variability as well as higher-order moments. Consequently, both accumulation rate ν_{drift} and dispersion rate ν_{noise} change proportionally with sensory input, and the scaling property is naturally obtained.

However, it does not ensure that the correct shape will be obtained for the FPT density.

In this chapter, we primarily focus on the second aspect of the scaling property, i.e. the specific shape of the FPT distribution, which rely on less intuitive factors, and thus required a more detailed treatment. In order to understand how shorter-tailed distributions may be obtained, we further investigate three different aspects of the GE process which contributes to the shape of the distribution: *granularity*,

local bistability, and *far-from equilibrium* operation. For this we will compare the GE process to two other related random walks: a balanced Poisson process (BP) and the process obtained in the diffusion limit of the GE process, a Cox-Ingersoll-Ross process (CIR) [Cox et al., 1985, Tuckwell, 1988/2008].

4.1 Generalized Ehrenfest process

We have shown in Chapter 3 that the GE process is able, not only to reproduce the scaling property, but also to account for the characteristic, Gamma-like shape of observed reversal time densities. The perceptually relevant aspects of these observations are illustrated in **Fig. 4.1** and **Fig. 4.2**. As previously established, the correct shape for the FPT distribution can only be obtained in the DDR, not in the NDR. In contrast, the modified WP or OU processes produce scale-invariant FPT distributions, but did not produce the correct distribution shape for any input and threshold combination. In the modified OU process in particular, the FPT distribution in both DDRs and NDRs is an inverse Gaussian (or Wald distribution), with $\gamma_1 = 3c_v$, and only differ in terms of their mean FPT (see **Fig. 2.13**).

To single out which characteristics and conditions allow the GE process to reproduce the characteristic Gamma-like FPT distribution, we propose the study two other stochastic processes. Both processes are somewhat related to the GE, in that they share some of its key features. In particular, while both BP and CIR processes naturally satisfy the first aspect of the scaling property (invariance of the shape as the mean varies), neither of them reproduce the second aspect of the scaling property, i.e. correct shape of FPT distributions. The comparative analysis of their FPT properties will indeed provide important additional indications, and provide an intuitive account to this second and yet unexplained aspect of the scaling property.

First, we introduce a Balanced Poisson process (BP), a discrete birth-death process which also accumulates activity with discrete Poisson events, but is not limited in size. Second, we introduce a Cox-Ingersoll-Ross process (CIR), a Gaussian diffusion model which shares the exact same drift and dispersion rates as the GE,

but is not discrete.

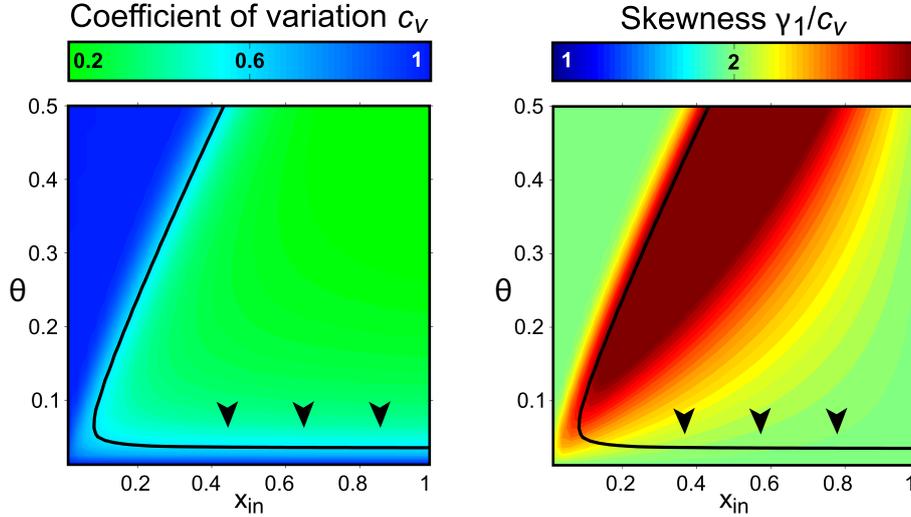


Figure 4.1: Generalised Ehrenfest process with $N = 80$ bistable units: normalized moments of FPT density as functions of input x_{in} and threshold θ (both in units of σ). (Left): coefficient of variation c_v . (Right): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$, black arrows mark several locations where a gamma distribution is obtained for different average dominance times. Reprinted from [Cao et al., 2016] with the co-authors' permission.

4.2 Balanced Poisson process

4.2.1 Definition and Master Equation

A (BP) process combines two independent Poisson processes contributing, respectively, activity increments and decrements of unit size with homogeneous rates ν_E and ν_I [Tuckwell, 1988/2008]. This process approximates excitatory and inhibitory post-synaptic potentials incrementing and decrementing the membrane potential of a neuron [Tuckwell, 1988/2008]. Excitatory (respectively inhibitory) spikes arrive with a fixed rate ν_E (respectively ν_I), so that either spike counts form independent Poisson processes $N_E(t)$ (respectively $N_I(t)$). We consider the net activation at time t , $N(t) = N_E(t) - N_I(t)$, as well as its distribution $P_n(t)$ (i.e. the probability of having a net activity $N(t) = n$ at time t). Within a small interval dt , we can

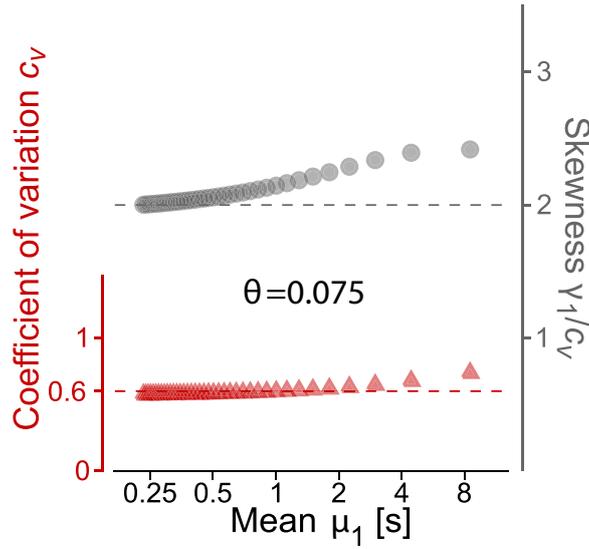


Figure 4.2: Generalised Ehrenfest process: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). For constant θ (value given by inset), μ_1 decreases as x_{in} increases. The scalar property is satisfied for both the CV and the Skewness, while the ratio γ_1/c_v 2 matches behavioural data. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

expect an increment with probability $\nu_E dt$, and a decrement with probability $\nu_I dt$, or none with probability $(1 - \nu_E dt)(1 - \nu_I dt)$. Therefore, the probability of having a net activity n at a later time $t + dt$ is:

$$P_n(t + dt) = P_{n-1}(t)\nu_E dt + P_{n+1}(t)\nu_I dt - (1 - \nu_E dt)(1 - \nu_I dt)P_n(t) \quad (4.1)$$

This yields the following differential equation:

$$\frac{dP_n(t)}{dt} = \nu_E P_{n-1}(t) + \nu_I P_{n+1}(t) - (\nu_E + \nu_I)P_n(t) \quad (4.2)$$

Expression (4.2) is the Master Equation for the BP process. As previously done for the GE process, we obtain the drift and dispersion rates (ν_{drift} and ν_{noise} , respectively) by solving differential equations for $\langle n \rangle$ and $\langle n^2 \rangle$. At first-order, we

get:

$$\begin{aligned} \frac{d}{dt} \left(\sum_n n P_n(t) \right) &= \frac{d\langle n \rangle}{dt} = \nu_E - \nu_I \\ \Rightarrow \langle n \rangle &= (\nu_E - \nu_I)t \end{aligned} \quad (4.3)$$

At second order, we get:

$$\begin{aligned} \frac{d}{dt} \left(\sum_n n^2 P_n(t) \right) &= \frac{d\langle n^2 \rangle}{dt} = 2(\nu_E - \nu_I)\langle n \rangle + \nu_E + \nu_I \\ \Rightarrow \langle n^2 \rangle - \langle n \rangle^2 &= (\nu_E + \nu_I)t \end{aligned} \quad (4.4)$$

Thus, drift and dispersion rates are:

$$\nu_{drift}^* = \nu_E - \nu_I \quad (4.5)$$

$$\nu_{noise}^* = \nu_E + \nu_I \quad (4.6)$$

4.2.2 Exact solution

The analytical solution to equation (4.2) can be obtained using the generating function, as defined in (3.16). In the case of the GE process, transition rates depend on the number n of active switches in the pool, and the Master Equation becomes a partial differential equation which must be solved by the method of characteristics. Here, transition rates do not depend on the net activation n of the system. The resulting differential equation for the generating function $G(z, t)$ is slightly simpler than for the GE process, and can be solved by separation of variables.

$$\tau \sum_n \frac{dP_n(t)}{dt} z^n = \tau \frac{\partial G(z, t)}{\partial t} = \left(\lambda z + \frac{\mu}{z} - 1 \right) G(z, t) \quad (4.7)$$

Where we have defined the characteristic time-constant $\tau = (\nu_E + \nu_I)^{-1}$ as well as the upward transition probability $\lambda = \nu_E \tau$ and the downward transition probability $\mu = \nu_I \tau$, so that $\lambda + \mu = 1$. We then solve the differential equation with initial

condition $P_0(0) = 1$, so that $G(z, 0) = 1$:

$$G(z, t) = e^{-\frac{t}{\tau}} e^{\frac{t}{\tau}(\lambda z + \frac{\mu}{z})} \quad (4.8)$$

Equation (4.8) can then be expanded in powers of z , so that the probability $P_n(t)$ can be identified. The solution is given by the following relation, where $I_n(x)$ is the modified Bessel function of the first kind (see Appendix A for details):

$$P_n(t) = \left(\frac{\lambda}{\mu}\right)^{\frac{n}{2}} e^{-\frac{t}{\tau}} I_n\left(2\frac{t}{\tau}\sqrt{\lambda\mu}\right) \quad (4.9)$$

4.2.3 First-passage-time distribution

The distribution of FPT T to a threshold θ from initial condition $N(0) = 0$, $f_\theta(T)$, can also be expressed in terms of the modified Bessel function of the first kind [Tuckwell, 1988/2008]:

$$f_\theta(t) = \theta \left(\frac{\nu_E}{\nu_I}\right)^{\theta/2} \frac{e^{-(\nu_E + \nu_I)t}}{t} I_\theta(2\sqrt{\nu_E \nu_I} t) \quad (4.10)$$

The two first moments of the distribution are also known [Tuckwell, 1988/2008]. Alternatively, the CV can also be obtained from (2.18), using (4.5) and (4.6). For the third raw moment, we made use of the following tabulated integral (see [Gradshteyn and Ryzhik, 1988], p700):

$$\int_0^\infty x^{m+1} e^{-\alpha x} I_\theta(\beta x) dx = (-1)^{m+1} \beta^{-\theta} \frac{d^{m+1}}{d\alpha^{m+1}} \left[\frac{(\alpha - \sqrt{\alpha^2 - \beta^2})^\theta}{\sqrt{\alpha^2 - \beta^2}} \right] ; m \geq 0 \quad (4.11)$$

We obtain:

$$\langle T^3 \rangle = \left(\frac{\theta}{\nu_E - \nu_I}\right)^3 + 3\theta^2 \frac{\nu_E + \nu_I}{(\nu_E - \nu_I)^4} + 3\theta \frac{(\nu_E + \nu_I)^2}{(\nu_E - \nu_I)^5}$$

So that:

$$\mu_1 = \frac{\theta}{\nu_E - \nu_I} \quad (4.12)$$

$$c_v = \sqrt{\frac{\nu_E + \nu_I}{\theta(\nu_E - \nu_I)}} \quad (4.13)$$

$$\gamma_1 = 3c_v \quad (4.14)$$

In this case, the stimulus strength is $x_{in} = (\nu_E - \nu_I)\tau = \lambda - \mu$. As previously, we consider the moments of the FPT density with respect to the threshold θ and the input-level x_{in} , as illustrated in **Fig. 4.3**. The FPT statistics of the BP process is similar to that of the modified WP with drift and input-dependent noise.

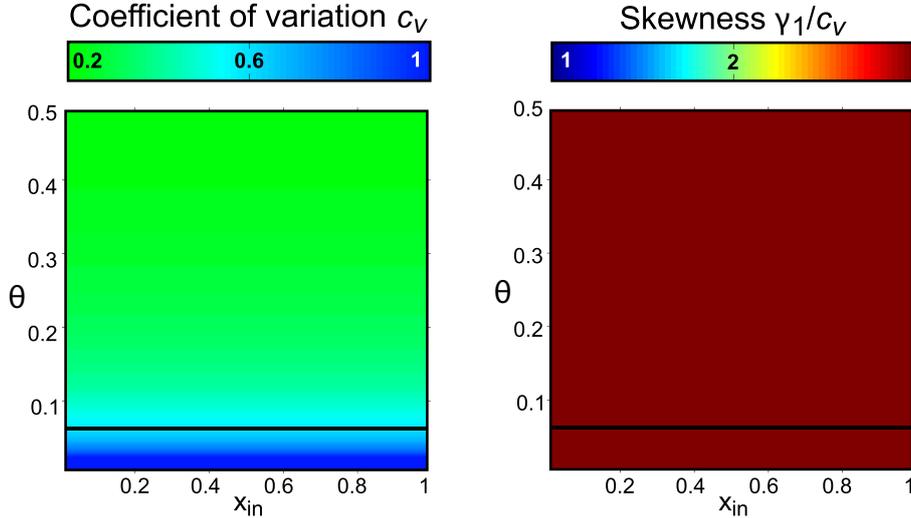


Figure 4.3: Balanced Poisson (BP) process: normalized moments of FPT density as functions of input x_{in} and threshold θ (both in units of σ). (*Left*): coefficient of variation c_v . (*Right*): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

Like the GE process, the dynamics of the BP process is composed of discrete spontaneous events whose rates depend on the input-level. Similarly, the event rate affects both the drift and the dispersion of the process, which is why the stochastic component of the dynamics naturally depends on the input. A scaling property is guaranteed (**Fig. 4.4**), provided sensory input s changes the event rates in a proportional manner, so that the proportionality between accumulation and dispersion is maintained, as required by (2.22). Again, the skewness $\gamma_1 = 3c_v$ is

larger than in experimental observations. This further confirms that the appropriate input-dependence of the noise is necessary and sufficient for the scaling property, but not for the distribution shape (we obtain an inverse Gaussian instead of a Gamma distribution).

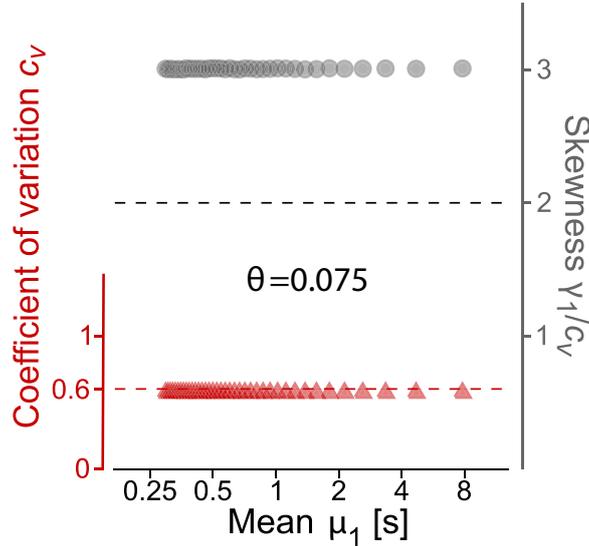


Figure 4.4: BP process: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). For constant θ (value given by inset), μ_1 decreases as x_{in} increases. The scalar property is satisfied for both the CV and the Skewness, but the ratio γ_1/c_v 3 is too large. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

4.3 Cox-Ingersoll-Ross process

4.3.1 System-size expansion

In this section, we derive a continuous equivalent of master equation (3.2) in the limit of large number N of switches in the pool. For this, we use the system-size expansion method [Hanggi et al., 1984]. The overall idea is that for a system of sufficiently large size, the ensemble of accessible states forms a continuum: the discrete nature of transitions between neighbouring states can be neglected, and the system can be described in terms of a continuous stochastic process.

We introduce the activity x (or fraction of active nodes in the pool):

$$x = \frac{n}{N} \quad (4.15)$$

In our case, the size parameter is the total number of nodes N , and we wish to derive a dynamical equation for the density $p(x, t)$, i.e. the probability to have a fraction x of active nodes at time t in the pool, from the the density $P(n, t)$.

The change of variable $n \rightarrow x$ requires these densities to scale as follows:

$$p(x, t) = NP(n, t) \quad (4.16)$$

Also, we must scale the original transition rates between neighbouring states:

$$W^\pm(n) = N\gamma^\pm(x) \quad (4.17)$$

We can now rewrite the original master equation (3.2) as a Fokker-Planck equation by expanding over small parameter $\epsilon = 1/N$:

$$\begin{aligned} \frac{1}{N} \frac{\partial p(x, t)}{\partial t} &= [1 - (x - \epsilon)]\nu_+ p(x - \epsilon) + (x + \epsilon)\nu_- p(x + \epsilon) - [(1 - x)\nu_+ + x\nu_-]p(x) \\ &= \gamma^+(x - \epsilon)p(x - \epsilon) + \gamma^-(x + \epsilon)p(x + \epsilon) - \gamma^+(x)p(x) - \gamma^-(x)p(x) \\ &= -\frac{1}{N} \frac{\partial}{\partial x} [(\gamma^+(x) - \gamma^-(x))p(x)] + \frac{1}{2N^2} \frac{\partial^2}{\partial x^2} [(\gamma^+(x) + \gamma^-(x))p(x)] \\ \tau \frac{\partial p(x, t)}{\partial t} &= -\frac{\partial}{\partial x} [(\lambda - x)p(x)] + \frac{1}{2} \frac{\partial^2}{\partial x^2} \left[\frac{\lambda(1 - x) + \mu x}{N} p(x) \right] \end{aligned} \quad (4.18)$$

Recall that $\lambda = x_{in}(s)$ indirectly represents the sensory input s , we identify the drift and diffusion terms:

$$\mu(x, x_{in}) = x_{in} - x \quad \sigma(x, x_{in}) = \sqrt{\frac{x_{in} - (2x_{in} - 1)x}{N}} \quad (4.19)$$

One can also be rewrite the partial differential equation (4.18), as a stochastic differential equation on variable $x(t)$, also known as the Langevin equation, where we define $\xi(t)$, an unitary Gaussian white noise, i.e. normally distributed with mean

zero and variance 1:

$$\tau \frac{dx}{dt} = \lambda - x + \sqrt{\tau} \sqrt{\frac{\lambda(1-x) + \mu x}{N}} \xi(t) \quad (4.20)$$

For both Eq.(4.18) and (4.20), taking the limit $N \rightarrow \infty$ causes the diffusion term to vanish. The dynamics is then reduced to its deterministic component, and we recover the first-order approximation previously obtained in (3.7).

4.3.2 Scaling property in the CIR process

Interestingly, the reduction of the Master Equation (3.2) for the GE process, to a Gaussian-diffusion type equation (expressions (4.18) and (4.20)) reveals a specific dependence of the noise on the input-level and the state of the system. Recall that here, the input corresponds to the asymptotic value of the process (i.e. the value for which the drift vanishes), so that $x_{in} = \lambda$. We can rewrite the infinitesimal variance of the CIR process as:

$$\sigma^2(x, x_{in}) = \frac{x_{in}}{N} - \frac{2x_{in} - 1}{N} x \quad (4.21)$$

Under this form, we see that the infinitesimal variance of the CIR process has two separate components.

The first component (first term on the right hand side of (4.21)) consist of an input-dependent additive noise. In fact, this dependence on x_{in} is the same we previously introduced in Chap. 2 (modified WP and OU processes). This correction allowed these models to satisfy the scaling property, in contrast to their unmodified, constant-noise counterparts. Since the GE process itself satisfies the scaling property, it is not surprising that its continuous approximation in the diffusion limit (the CIR process) would naturally exhibit this input-dependence of the noise amplitude.

The second component (first term on the right hand side of (4.21)) consist of a multiplicative noise term, both input- and state-dependent. The consequences of this specific structure of fluctuations on the scaling property, and on the shape of

the FPT distribution will be discussed in the next section.

Although it is clear that the CIR process will fulfil the scaling property, since its infinitesimal variance already satisfies the required input-dependence, we can derive the drift and dispersion rates ν_{drift} and ν_{noise} . The mean and variance for the CIR process are known [Cox et al., 1985]. Here we take the initial condition $x_0 = 0$.

$$\langle x \rangle = \lambda(1 - e^{-\frac{t}{\tau}}) \quad (4.22)$$

$$\langle x^2 \rangle - \langle x \rangle^2 = \frac{1}{N} \left[\lambda\mu + \lambda(\lambda - \mu)e^{-\frac{t}{\tau}} - \lambda^2 e^{-\frac{2t}{\tau}} \right] \quad (4.23)$$

Equations (4.22) and (4.23), which describe respectively the average dynamics of the CIR process, and its variance, are identical to (3.7) and (3.11). This is because the CIR process is simply a truncated version of the GE process, where fluctuations from third-order onwards are neglected in the system-size expansion. Up to the second order in fluctuations, both the GE process and the CIR process are identical in how they accumulate and disperse activity. Drift and dispersion rates are also the same (see (3.12)):

$$\begin{aligned} \nu_{drift} &= \frac{\lambda}{\tau} e^{-\frac{t}{\tau}} \\ \nu_{noise} &= \frac{\lambda}{\tau N} e^{-\frac{t}{\tau}} \left[2\lambda e^{-\frac{t}{\tau}} - (\lambda - \mu) \right] \end{aligned} \quad (4.24)$$

And in the limit of a DDR, we get:

$$\begin{aligned} \nu_{drift}^* &= \frac{\lambda}{\tau} \\ \nu_{noise}^* &= \frac{\lambda}{\tau N} \Rightarrow c_v^* = \frac{1}{\sqrt{\theta N}} \end{aligned} \quad (4.25)$$

The CIR process shows similar FPT statistics than the discrete GE model from which it is derived (see **Fig. 4.5**). Isoclines of the CV also reveal two distinct regimens in which the condition $c_v \approx 0.6$ can be obtained. As before, we will focus on the DDR, where the threshold can be held constant, as opposed to the NDR,

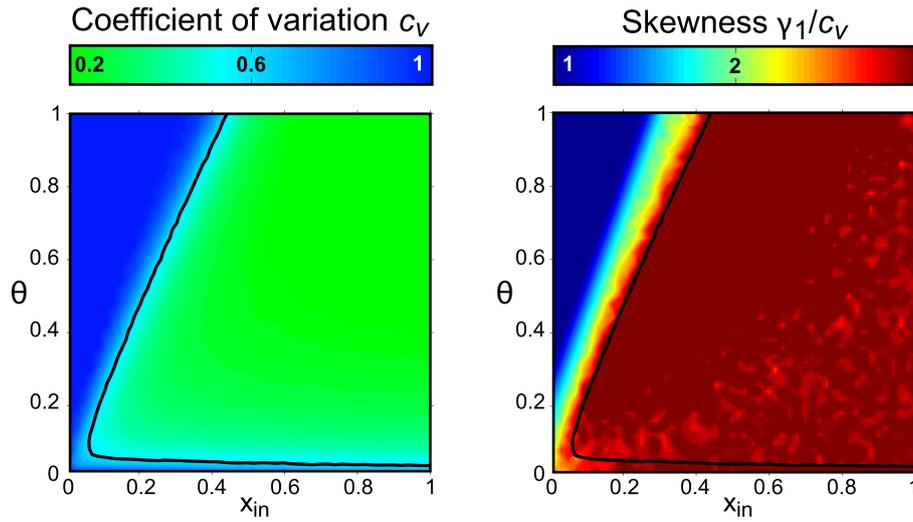


Figure 4.5: CIR process: normalized moments of FPT density as functions of input x_{in} and threshold θ (both in units of σ). (Left): coefficient of variation c_v . (Right): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$. Reprinted from [Cao et al., 2016] with the co-authors' permission.

where the value of the threshold must be continuously adapted to the input-level. Again, in a DDR, the value of the threshold must be kept low to obtain sufficiently variable dominance durations.

Within these conditions, the scaling property is observed (see **Fig. 4.6**). However, unlike the GE process, the shape of the distribution corresponds to a inverse Gaussian with $\gamma_1 \approx 3c_v$. Since the CIR process is a Gaussian approximation of the GE process, we deduce that differences in their FPT statistics correspond to the non-Gaussian aspect of the latter. In the NDR, along the $c_v = 0.6$ isocline, both processes are equivalent and share qualitatively similar FPT statistics, indicating that the GE is equivalent to a Gaussian process around steady-state. This is not the case in the DDR, which reveals one of the key dynamical features uniquely expressed by the GE process: when operated sufficiently far from equilibrium, it diffuses *asymmetrically*.

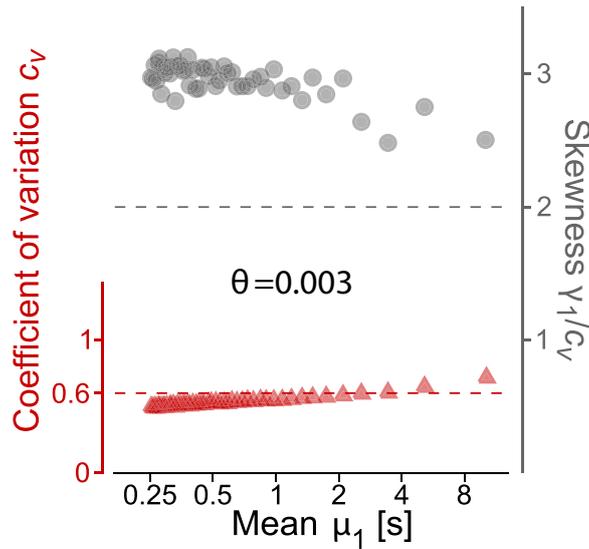


Figure 4.6: CIR process: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). For constant θ (value given by inset), μ_1 decreases as x_{in} increases. The scalar property is satisfied for both the CV and the Skewness, but the ratio γ_1/c_v 3 is too large. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

4.4 Distinguishing properties of the GE process

4.4.1 Overview

So far we have studied the FPT statistics of a number of random models. These results are summarized in Table 4.4.1. This should highlight two important points. Firstly, that the scaling property requires a specific balance between deterministic and stochastic components of the dynamics to be observed and maintained, while input modulations may cause large deviations of mean FPTs. Specifically, we have shown in Chapter 2 that such modulations of the input-level must affect the drift and dispersion rates (respectively ν_{drift}^* and ν_{noise}^*) in a proportional manner. Processes such as the WP or OU processes, with constant noise amplitude, do not satisfy the scaling property, and can be modified by assuming a certain input-dependence of the noise amplitude, so that the scaling property is obtained. When accumulating activity with bistable nodes, the input-dependence of events rates ensures that drift and noise rates will vary in a proportional manner, as is

the case for the BP process and the GE process. This is also highlighted by the fact that the CIR process, which is obtained by taking the continuous limit of the GE process, also inherits from the same input-dependence of the noise amplitude as the modified WP and OU processes.

Name	Type	ν_{drift}^*	ν_{noise}^*	Scaling prop.	$c_v \approx 0.6$ & $\gamma_1 \approx 2c_v$
WP	Gaussian	$\frac{x_{in}}{\tau}$	$\frac{\sigma^2}{\tau}$	–	–
OU	Gaussian	$\frac{x_{in}}{\tau}$	$\frac{\sigma^2}{\tau}$	–	–
Mod. WP	Gaussian	$\frac{x_{in}}{\tau}$	$\frac{x_{in}}{\tau\alpha}$	✓	–
Mod. OU	Gaussian	$\frac{x_{in}}{\tau}$	$\frac{x_{in}}{\tau\alpha}$	✓	–
CIR	Gaussian	ν_+	$\frac{\nu_+}{N}$	✓	–
BP	Infinite Birth-Death	$\nu_E - \nu_I$	$\nu_E + \nu_I$	✓	–
GE	Finite Birth-Death	ν_+	$\frac{\nu_+}{N}$	✓	✓

Table 1 Comparison of investigated random-walk processes, ordered in three sections (separated by double lines). Processes with *constant noise amplitude* do not satisfy the scaling property (WP and OU). Processes with *input-dependent noise amplitude* (modified WP, modified OU, CIR and BP): normalized moments remain constant with input and a scaling property is obtained (in the DDR), because accumulation rate ν_{drift}^* and dispersion rate ν_{noise}^* increase with input in the same proportions. The scaling property, as well as the distribution shape, is uniquely obtained by a finite birth-death process, the GE process.

Secondly, that the characteristic, Gamma-like shape of observed reversal times densities, which implies that the skewness obeys the relation $\gamma_1 \approx 2c_v$, can only be obtained by the GE process. We propose that the success of this model in reproducing the observed statistics of multistable perception, including the characteristic Gamma-like shape, relies on several key features:

- *Granularity*: Stochastic accumulation of activity is performed collectively by

a *finite* number N of microscopic units

- *Local bistability*: Microscopic states are locally bistable units with spontaneous activation and inactivation rates (ν_+ and ν_- , respectively)
- *Input-dependence*: The relative rate ν_+/ν_- , and thus also the steady-state activation $x_{in} = \nu_+ / (\nu_- + \nu_+)$, grows with sensory input s .
- *Far-from-equilibrium regime*: Both initial and threshold activation are comparatively small and far from steady-state, $x_0 < \theta \ll x_{in}$

4.4.2 Granularity, finite-size effects

In the GE process, changes in activity are produced by discrete increments and decrements, which originate within a finite population of spontaneously bistable units. Note that in the Master Equation (3.2) transition rates between activity states depend on the number n of active units. Upward transition rates, given n active units, depend on the number $N - n$ of inactive units, which may undergo an upward transition, so that the total increment rate is: $W_n^+ = (N - n)\nu_+ = N\nu_+(1 - x)$. If n units are active, each of these units may undergo a downward transition, so that the total decrement rate is: $W_n^- = n\nu_- = N\nu_-x$. This means that if the population depleted, i.e. fully inactive, only upward transitions may occur. Conversely, if the population is saturated, i.e. fully active, only downward transitions may occur. This means that whenever the population activity is close to zero and accumulates towards a comparatively large equilibrium value x_{in} , negative accumulation caused by active units switching downwards will be negligible (see **Fig. 4.7**)

The drift term $\mu(x, x_{in})$, as well the dispersion term $\sigma(x, x_{in})$, in the Fokker-Planck description of the CIR process, both depend on the activity and input-level. As mentioned previously, the variance $\sigma^2(x, x_{in})$ is composed of an additive, input-dependent component, and of a multiplicative, input-dependent component (see (4.21)).

A direct consequence of this particular noise structure, is that the variance of

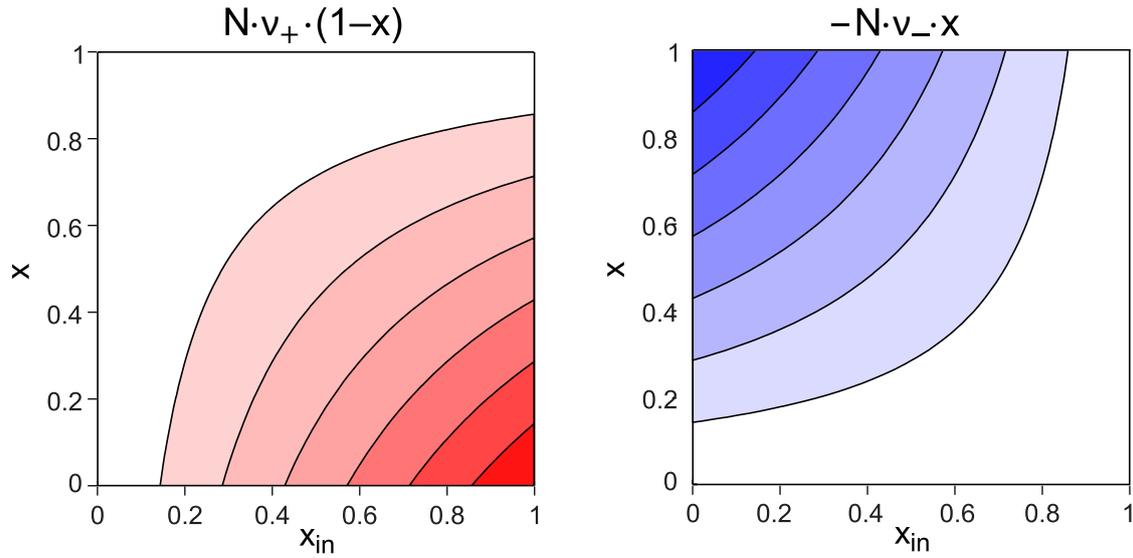


Figure 4.7: Increment and decrement rates depend on the input-level x_{in} and the activity x . (*Left*): Increment rate, counted positively (red) increases with x_{in} and decreases with x . Large increment rates are obtained in the region $x \ll x_{in}$, where decrement rates are negligible. (*Right*): decrement rate, counted negatively (blue) increases with x_{in} and decreases with x . Large decrement rates are obtained in the region $x \gg x_{in}$, where increment rates are negligible. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

the process $\sigma^2(x, x_{in})$, unlike the drift $\mu(x, x_{in})$, is non-monotonic over the whole (x, x_{in}) -space. For instance, if $x < 0.5$, $\sigma^2(x, x_{in})$ increases for increasing x_{in} , but, if $x > 0.5$, $\sigma^2(x, x_{in})$ decreases for increasing x_{in} . We know that to satisfy the scaling property, the variance of the process should increase for increasing input-level (as explained in **Fig. 2.16**), therefore, the multiplicative noise term implies that $x < \theta < 0.5$.

Furthermore, the scaling property requires that variations of the drift $\mu(x, x_{in})$ caused by variations of x_{in} or x must be appropriately matched by the variations of $\sigma^2(x, x_{in})$, so that the condition $\nu_{drift} \propto \nu_{noise}$ may hold. This effectively imposes boundaries on the range of value that the activity will take from the initial condition x_0 to the threshold θ (see **Fig. 4.8**).

When assuming an initial condition so that $x_0 < \theta < 0.5$, but where x_0 is significantly greater than zero, both the drift and variance of the process increase with the input x_{in} , and with the activity x as it accumulates towards the threshold.

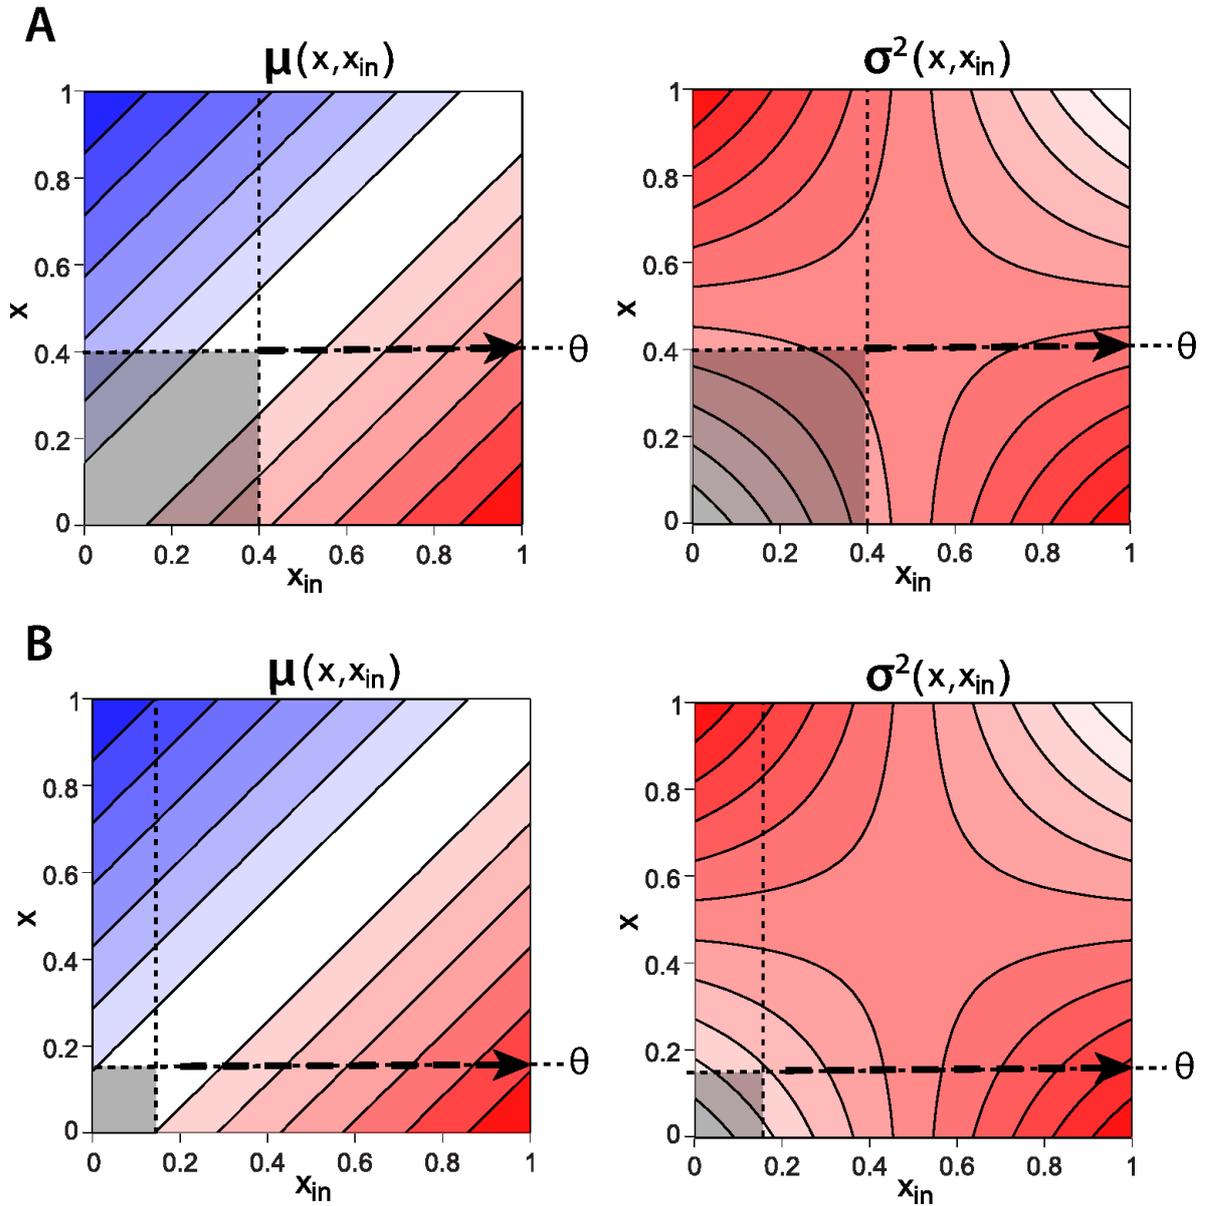


Figure 4.8: Drift $\mu(x, x_{in})$ and variance $\sigma^2(x, x_{in})$ of the GE process as a function of the activity x and input-level x_{in} . Dotted lines mark the position of the initial activity and threshold. Shaded areas correspond to the NDR ($\theta > x_{in}$) **A** (Left): when the initial activity x_0 is significantly greater than zero, the drift increases with input level x_{in} (as the difference between total increment and decrement rates, see **Fig. 4.7**) and as activity x grows from x_0 to θ (dashed arrow). (Right): The variance increases as the sum between total increment and decrement rates, but comparatively less (dashed arrow). **B** For $x \ll 0.5$, both drift (left) and dispersion (right) increase in a proportional manner (dashed arrows), ensuring that normalized moments remain constant. Adapted from [Cao et al., 2016] with the co-authors' permission.

However, the increase in the drift term is comparatively greater than the increase of the dispersion term (**Fig. 4.8 A**). When assuming $x_0 < \theta \ll 0.5$, so that the initial activity is close to zero, the increase in both drift and dispersion, as both the input level and the activity increase, are comparable (see **Fig. 4.8 B**). Therefore, in order to satisfy the scaling property, i.e. to maintain the balance between the deterministic component of the dynamics, not only the range $[x_0, \theta]$ must be small, but also, the initial value x_0 must be close enough to zero. This ensures a sufficiently strong input-dependence of the dispersion of the process, with respect to the drift. This ties in with the conclusion of the regime identification performed in Chapter 3 (see **Fig. 3.7**): both the scaling property, and the distribution shape (with $\gamma_1 \approx 2c_v$) were best verified when considering a strong DDR, with an initial condition close to zero.

4.4.3 Asymmetric dispersion far-from-equilibrium

The CIR process, and the GE process, from which the former is derived, can be considered identical up to the second-order in fluctuations. In other words, their average dynamics, as well as the dynamics of their variance, is the same (see (3.7), (3.11), (4.22), and (4.23)). Both models obey the scaling property under the same assumptions, while only the GE model can reproduce FPT densities with the same skewness as those observed experimentally. Note that the same observation is also valid for the modified OU process introduced in Chapter 2. All three processes have the same FPT distributions in the noise-dominated regime, corresponding to an inverse Gaussian distribution (see **Fig. 2.13**, **Fig. 4.5** and **Fig. 4.1**), and differ in the drift-dominated regime, where only the GE has a Gamma-like FPT distribution.

To understand why this is so, one must also consider fluctuations up to the third-order (skewness). Indeed, both modified OU and CIR processes assume a Gaussian random walk dynamics with normally distributed steps; thus, by definition, the probability density of the process, at any given time, is symmetric. In the GE process, the temporal evolution of the probability density corresponds to a binomial distribution, with a time-varying parameter $p = \lambda(1 - e^{-\frac{t}{\tau}})$ (see (3.21)). The bino-

mial distribution is often well approximated by a normal distribution (specifically, $\mathcal{N}(np, np(p-1))$, where n is the number of realisation of the Bernoulli random variable), and therefore, one could think that the GE process also diffuses symmetrically, in the manner of a Gaussian random walk.

The skewness of a binomial distribution of parameter p is given by the formula:

$$\gamma_1^{binom.} = \frac{1-2p}{\sqrt{np(1-p)}} \quad (4.26)$$

For a sufficiently large number of realisations n , the above expression should tend to zero. However, this is not the case in the situations we are considering, where the initial activity is very close to zero. In this case, the binomial distribution is right-skewed. Moreover, our process accumulates activity up to a low-threshold, in a strong DDR, which implies $t \ll \tau$, and thus the parameter p remains close enough to zero for the distribution to remain skewed.

The asymmetric nature of the diffusion in the vicinity of the boundary can be qualitatively understood by considering the case where the assembly of bistable units is fully depleted, and accumulates activity towards a sufficiently distant steady-state. Over time, the distribution of activity is expected to shift in the direction of the steady-state under the influence of the drift (or deterministic component of the dynamics). However, the total decrement rate, which is proportional to the number of active units, is initially equal to zero, simply because the population is completely inactive: there is no bistable unit left to deactivate. For this reason, the stochastic component of the dynamics, can only disperse activity in the same direction as the drift, resulting in an asymmetric diffusion (**Fig. 4.9 A**).

The same reasoning hold as long as the initial value of activity is close enough to zero so that positive accumulation is strong and negative accumulation negligible (see **Fig. 4.7**). Such non-Gaussian, asymmetric dispersion, thins out realizations lingering near the initial value, keeping the activity distribution compact and asymmetric. The comparatively compact dispersion of activity then translates into

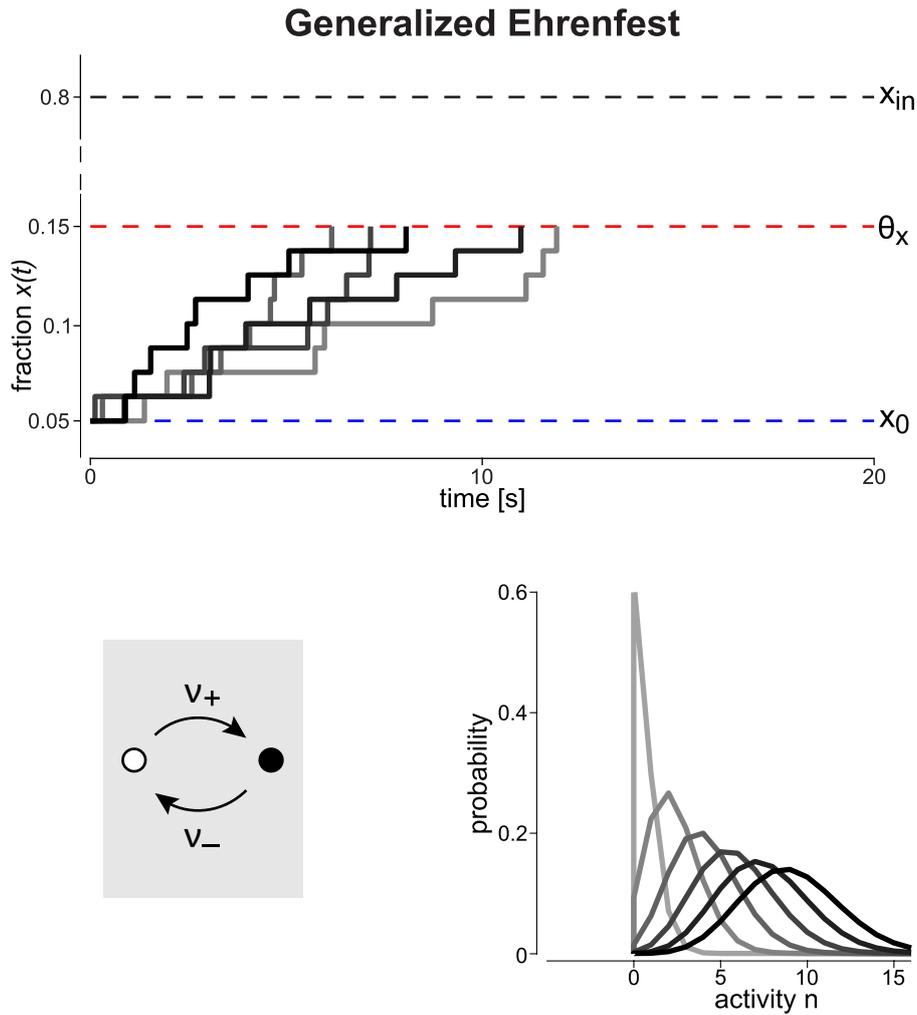


Figure 4.9: GE process of fractional activity $x(t)$ in a population of $N = 80$ units. Units are bistable and activate or inactivate spontaneously with rates $\nu_+ = 0.008 \text{ Hz}$ and $\nu_- = 0.002 \text{ Hz}$, respectively. (Top): Accumulation begins at $x_0 = 0.05$ and ends at $\theta_x = 0.15$. Note that activity remains well below the steady-state activity level $x_{in} = \nu_+ / (\nu_+ + \nu_-)$. (Bottom-right): Distribution of activity $x(t)$ at different times t , with darker hues representing later times. Asymmetric dispersion reflects the constraint $x_0 = 0$. Reprinted from [Cao et al., 2016] with the co-authors' permission.

a short-tailed FPT distribution with skewness $\gamma_1 \approx 2c_v$. These late realizations are precisely those which, if the diffusion was symmetric, would accumulate in the tail of the FPT distribution, thereby increasing its skewness, as has been the case for all other models scrutinized in this study but the GE process. This boundary effect is also apparent in **Fig. 4.1 (Right)**: the appropriate distribution, with $\gamma_1 \approx 2c_v$, becomes increasingly difficult to obtain for a wide range of average FPT

if one assumes larger initial condition for the activity.

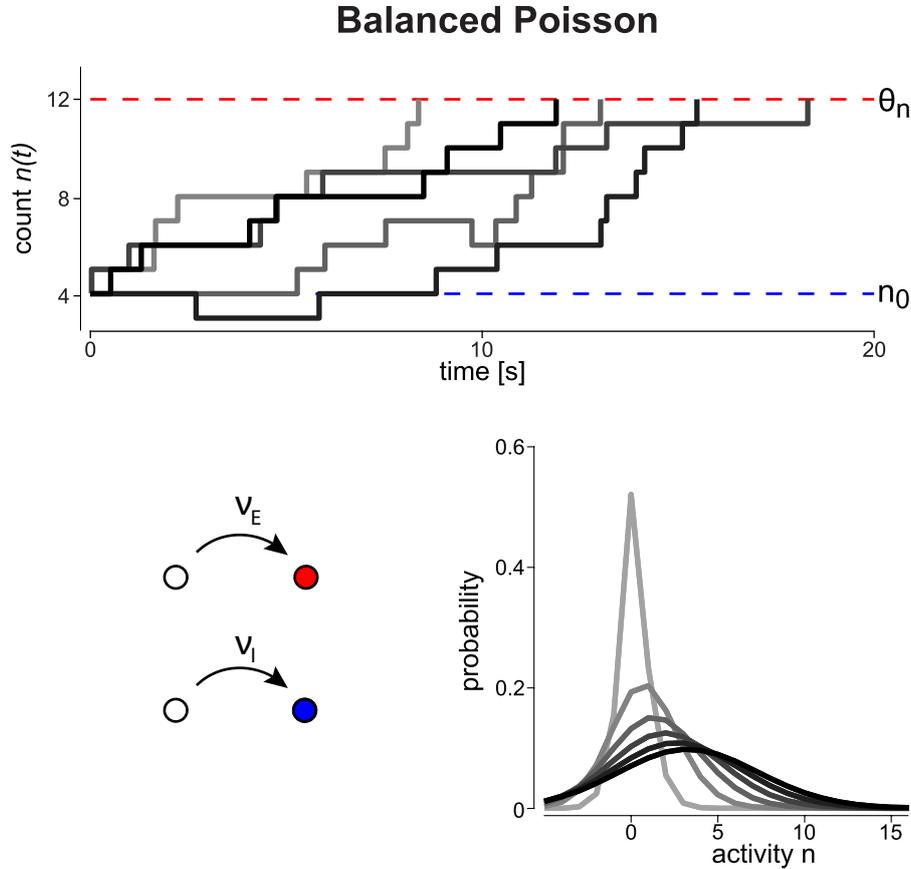


Figure 4.10: (Top): BP process counting independent increments and decrements with rates $\nu_E = 0.64 \text{ Hz}$ and $\nu_I = 0.16 \text{ Hz}$, respectively. (Top): Accumulation begins at $n_0 = 4$ and ends at $\theta_n = 12$. Five realizations are shown (different shadings). (Bottom-right): Distribution of activity $x(t)$ at different times t , with darker hues representing later times. Here, dispersion is symmetric, allowing low activity (including negative counts) to remain probable. Reprinted from [Cao et al., 2016] with the co-authors' permission.

In comparison, the BP process, which also accumulates activity in a sequential manner - in the sense that increments and decrements of activity follow from spontaneous and discrete events - and thus satisfies the scaling property, disperses activity in symmetrically and produces FPT distributions with heavier-tails than those observed experimentally (**Fig. 4.9 B**). The qualitative difference between the GE and Poisson process is somewhat more salient when observing the evolution over time of their respective probability densities $[p_n(t)]$ (compare **Fig. 4.9** and **Fig. 4.10**). There, we matched the increment and decrement rates of the Poisson process, so

that maximum increment and decrement rates in the GE process are identical. For the BP process, we set $\nu_E = 0.64\text{Hz}$ and $\nu_I = 0.16\text{Hz}$. For the GE process, with $N = 80$ bistable units, the maximum increment rate is obtained assuming all units are inactive and the maximum decrement rate is obtained assuming all units are active, which yields $\nu_+ = 0.008\text{Hz}$ and $\nu_- = 0.002\text{Hz}$.

This is because unlike the GE process, increments and decrements are generated by a counting process which does not assume any boundary. There is no saturation effect, because there is no upper limit on the number of Poisson events producing increments of activity with rate ν_E . Similarly, there is no depletion effect, because there is no lower limit on the number of Poisson events producing decrements of activity with rate ν_I : compared to a GE process, more realisations are allowed to diffuse around the origin, and ultimately accumulate in the tail of the FPT distribution, yielding a greater skewness.

4.5 Neurally plausible realization of a GE process

As previously mentioned, our postulated bistable units are thought to idealize the dynamics of local attractors assemblies. They constitute an abstract model for the dynamics of recurrently connected neural structures, such as cortical columns. Qualitatively similar assemblies have already been implemented in networks of spiking-neurons, with recurrent excitation allowing the formation of self-sustaining high-activity states, and with transitions between states modulated by the level of synaptic activity [Amit, 1995, Amit and Brunel, 1997].

The GE process could be thought of as a collection of such assemblies. Note that we have shown that the GE process conserves its relevant dynamical properties when considering interactions between bistable units (see 3.4.3). The resulting structure forms a multi-modular, or clustered neural network, composed of either independent or coupled recurrent attractor subnetworks (similarly to [Litwin-Kumar and Doiron, 2012]). We now propose a spiking-neuron implementation of the GE process and study how it collectively accumulates activity, compared to a non-clustered network, i.e. which is not formed of a collection of locally bistable

subnetworks.

The results of this simulation are illustrated in **Fig. 4.11 A**. Each such assembly was composed of *recurrently and strongly coupled* spiking neurons. Bistable attractor dynamics was achieved by balancing excitatory and inhibitory projections within each of these assemblies. In total, 20 sub-networks were simulated, with each comprising 125 excitatory leaky-integrate-and-fire neurons ('foreground'), which were weakly coupled to an additional 875 excitatory neurons ('background'), as well as 250 inhibitory neurons. Overall connection probability between any two neurons was $c = 80\%$.

Spontaneous activations and inactivations in 'foreground' activity between a low state of approximately 3 Hz and a high state of approximately 40 Hz occurred, driven by endogenous activity fluctuations (finite-size noise). For each of the 20 bistable units, local switching rates changed with external input, typically increasing for increasing input-levels.

Collective stochastic accumulation of activity from a near quiescent state, to a near saturated state was obtained, among the population of bistable units, by *abruptly* altering the level of external stimuli. This caused the overall steady-state in the population to raise from a near-quiescent state (fraction of active units close to zero), to a fully-active state (fraction of active units close to one), with corresponding low and high states of activity. The spike raster in **Fig. 4.11 A** shows individual assemblies successively activating at different (random) times and, superimposed, the gradual accumulation of average activity in the network (mean instantaneous firing rate over all assemblies, noted $r(t)$ in the figures). Average collective activity accumulates according to a predictable exponential profile, however, this simple behaviour of the average is deceptive. The detailed microscopic dynamics at the single-neuron level, as well as the detailed mesoscopic dynamics at the level of the ensemble of bistable assemblies is considerably more complex, as described from a more analytical perspective in Chapter 3. If the threshold for collective activity is low, so that the entire accumulation proceeds far

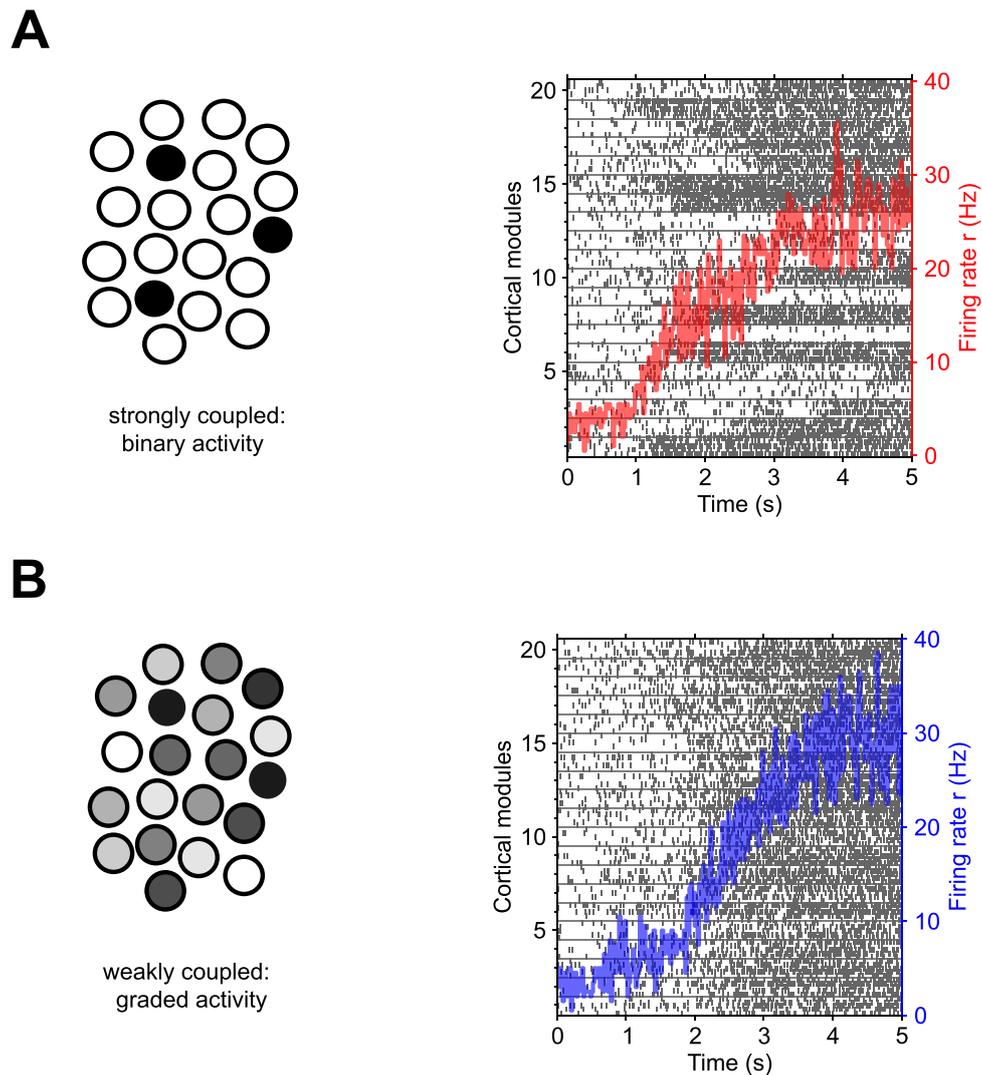


Figure 4.11: Stochastic accumulation of collective activity by modular assemblies of spiking neurons. **A:** strongly coupled assemblies expressing bistable attractor dynamics. Spike raster of five representative neurons per assembly (left ordinate) and collective activity $r(t)$ of all assemblies (red trace, right ordinate). Assemblies transition spontaneously and abruptly from inactive to active states (see text for details). **B:** weakly coupled assemblies without bistable dynamics. The activity of each assembly fluctuates about a steady-state level which rises progressively due to external input. Collective activity $r(t)$ is comparable to **A** (blue trace, right ordinate). *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

from steady-state, the first-passage times of collective activity will reproduce the scaling property of multi-stable perception.

For comparison, we also illustrate an alternative scenario which accumulates

collective activity in a comparable way in terms of the instantaneous activity $r(t)$ (**Fig. 4.11 B**). There, in contrast with the previous case, the network consisted of *weakly coupled* assemblies of spiking neurons, which *did not express local bistable attractor properties*. Rather than transitioning between two states of low and high firing rate, each assembly saw its activity fluctuating, also under the influence of endogenous, finite-size noise, around a single steady-state. The position of this steady-state was set by external input, so that collective stochastic accumulation from a low to a high activity state could be simulated by *gradually* increasing the external input. The spike raster shows individual assemblies becoming progressively more active and, superimposed, the collective activity from all assemblies. While the accumulation of activity in either scenario (**4.11AB**) was comparable in terms of $r(t)$, the scaling property could not be obtained in the case of weakly coupled assemblies, because of the absence of the local attractor property.

For strongly coupled assemblies, excitatory synaptic efficacy between ‘foreground’ neurons, ‘background’ neurons’, and between the two was $J_{\text{fore}} = 0.618mV$, $J_{\text{back}} = 0.438mV$, and $J_{\text{inter}} = 0.402mV$, respectively. Inhibitory synaptic efficacy was $J_I = -1.50mV$ and the efficacy of excitatory synapses onto inhibitory neurons was $J_{IE} = 0.560mV$. Finally, ‘foreground’ neurons, ‘background neurons’, and ‘inhibitory neurons’ each received independent Poisson spike trains of $2340Hz$, $2280Hz$, and $2280Hz$, respectively. Other settings were as in [Mattia et al., 2013]. Weakly coupled assemblies were obtained by reducing synaptic efficacies $J_{\text{fore}} = 0.566mV$ and $J_{\text{back}} = 0.431mV$, as well as increasing efficacies $J_{\text{inter}} = 0.409mV$ (to maintain overall level of activity). For ‘foreground’ neurons (each with external Poisson inputs of $2400Hz$) the firing rate was approximately $3Hz$. To reproduce the gradual accumulation of activity by strongly coupled assemblies, we increased (at a suitable pace) external Poisson inputs to $2760Hz$, eventually raising the firing rate to approximately $30Hz$.

By construction, the alternative scenarios illustrated by **Figs. 4.11 AB** are indistinguishable in terms of the time-varying average spike rate $r(t)$. Neverthe-

less, even at the single-cell level, a bistable dynamics of individual assemblies is detectable in terms of the variability of firing rates $\Delta r(t)$ and the distribution of inter-spike-intervals. To support this last statement, we provide a comparative analysis of the spiking activity of ‘foreground’ neurons with a population of non-homogeneous Poisson point processes, or ‘surrogate’ neurons. Statistics of spiking activity within the ‘surrogate’ population was generated to exactly match the instantaneous average firing rate $r(t)$ measured in the ‘foreground’ population. For strongly coupled assemblies expressing a local bistable attractor property, the standard deviation $\Delta r(t)$ of the instantaneous firing rate showed a significant increase during the rising phase of the accumulation, i.e. as soon as the abrupt change in external stimulus level occurred (**Figs. 4.12**, Left, red trace). For weakly coupled assemblies, the standard deviation of the instantaneous firing rate remained comparable to the one measured in the surrogate population of inhomogeneous Poisson spikes (**Figs. 4.12**, Left, blue trace). Additionally, we observed important differences in the distribution of interspike-intervals between spike trains obtained in both weakly and strongly coupled scenarios and the surrogate population. The inter-spike-interval distribution for the network constituted of locally bistable attractors emphasized extremes (short and long intervals) at the expense of the middle (**Figs. 4.12**, Right). This clearly reflects the local bistable dynamics of recurrently interacting cortical columns, with short ISI corresponding to the high firing-rate within active bistable units, and long ISI corresponding to the low firing-rate of the (few) bistable units which remained inactive despite the rising activity of the whole network. In weakly coupled units, the statistics of ISI mostly resembles that of a corresponding inhomogeneous Poisson with same average firing rate.

4.6 Summary

To summarize, recall that the scaling property presents two separate aspects: the invariance of shape of reversal time densities on the one hand, and the characteristic shape of reversal time densities, on the other hand. The first aspect has been discussed in details in Chapter 2 and 3, and can be ensured by considering

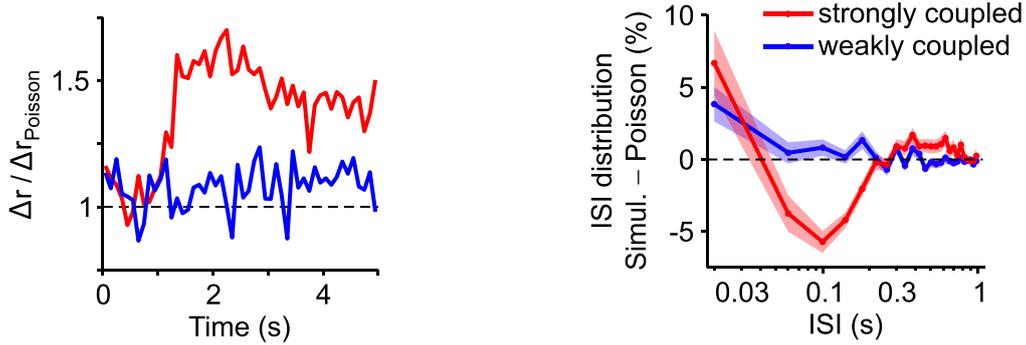


Figure 4.12: Comparison to inhomogeneous Poisson statistics. (*Left*): Standard deviation Δr of mean firing rate $r(t)$ in **Figs. 4.11 AB** (red and blue trace, respectively), computed in 100 ms sliding windows relative to the SD $\Delta r_{\text{Poisson}}$ of inhomogeneous Poisson processes reproducing the observed $r(t)$. The SD may be computed either over neurons or over trials. (*Right*): Distribution of inter-spike intervals (ISI) in **Figs. 4.11 AB** (red and blue trace, respectively), relative to the ISI distribution of inhomogeneous Poisson processes. Colored shading indicates the SD of ISI density. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

the accumulation of activity with *locally bistable* units, because both accumulation and dispersion rate change proportionally with sensory input.

In this chapter we have focused on the second aspect of the scaling property. Comparing the GE process to the BP and CIR process, both of which satisfy the first, but not the second aspect of the scaling property, allowed us to single out the key requirements allowing the GE process to uniquely satisfy both aspects. Specifically, the GE process must operate in a *far-from-equilibrium, low-threshold regime*, to ensure that collective accumulation of activity disperses *asymmetrically*. Furthermore, that asymmetric dispersion can only be obtained when considering discrete accumulation in a *finite* assembly of bistable units.

It is worth mentioning that two previous attempts to model multistable perception by considering stochastic accumulation of activity in finite populations of bistable units [Taylor and Aldridge, 1974, Gigante et al., 2009] failed to reproduce this statistics, because neither operated far enough from equilibrium. In the context of the timing of perceptual reversals, this asymmetric, non-Gaussian feature of the GE process accounts for the third-order statistics of dominance durations (the skewness, $\gamma_1 \approx 2c_v$).

The concept of non-diffusivity has recently been invoked on the basis that the higher-order properties of a variety of multistable biological processes may reflect correlated underlying stochastic processes [Freyer et al., 2009]. Descriptions of the Fokker-Planck or Langevin type, which rely on the diffusion approximation, assume uncorrelated fluctuations within the population they describe. They are more characteristic of systems fluctuating around equilibrium, and typically interacting on a short-range. This is at odds with many aspects of brain dynamics, which is typically out-of-equilibrium, distributed, and hierarchical. This chapter highlighted similar concepts, in that far-from-equilibrium operation, in a regime where fluctuations become non-Gaussian may be necessary to capture all aspects of perceptual dynamics.

Hierarchical Multistable Dynamics

In previous chapters, we have modelled the timing of perceptual reversals from the perspective of a first-passage-time problem, and used various diffusion-to-bound models to represent the accumulation of neural activity leading to a reversal. We identified the dynamical regimes in which some stochastic processes may be able to reproduce the empirical statistics of reversal times. Until now, in this admittedly oversimplified framework, we assumed perceptual reversals to occur automatically and did not seek to implement the mechanism driving perceptual reversals. These reductions were necessary in order to comprehensively account for important aspects of the fluctuating neural dynamics underlying multistable perception. This allowed us to identify which aspects of such stochastic dynamics may be at the origins of the puzzling scaling property observed in behavioural experiments. This extended preliminary study prepared the ground for an adequate implementation of a reversal mechanism, so that properties such as the scaling property would pertain in the resulting model.

Here, we show that the well-established and somewhat counter-intuitive experimental observations known as ‘Levelt’s propositions’, along with the scaling property, fully constrain the hierarchical structure and dynamics driving perceptual reversals. Our model quantitatively accounts for important aspects of behavioural observations, and uncovers simple qualitative mechanisms to explain Levelt’s propositions, the scaling property, and the characteristic Gamma shape of reversal time densities. Additionally, our hierarchical model predicts unexpected non-stationary properties

in alternation series. Preliminary results are proposed which confirm this predictions in experimental observations.

5.1 Introduction

The mechanisms underlying perceptual reversals constitute the core of an ongoing, and long-standing debate for which a definite solution has not yet been formulated. A common denominator to all proposed theories is that multistability arises from the structure and dynamics of interactions between the respective neural substrates representing alternative perceptual appearances. The extensive amount of observations from perceptual experiments have provided evidences for several processes thought to contribute to reversal dynamics, which have guided theoretical studies in designing computational models. Here, we will present the results which are most relevant to our proposal.

A first important point is the existence of a fast and a slow time-scale in perceptual dynamics: reversals are typically much shorter the typical time between reversals. Although some multistable displays, such are binocular rivalry, are more prone to indeterminate or mixed percepts (e.g. binocular fusion or piecemeal rivalry in BR), they typically occur with small predominance [Hollins, 1980, van Ee, 2009, Pastukhov and Klanke, 2016]. Perceptual appearances are mostly mutually-exclusive, so that only one, or the alternative is reported at a time. This indicates that perceptual competition is resolved categorically and rapidly. In models of multistable perception, such ‘winner-take-all’ behaviour is achieved by assuming strong mutual-inhibition between alternative representations of the stimulus [Lehky, 1988, Laing and Chow, 2002, Wilson, 2003, Seely and Chow, 2011]. Competition between the relevant neural substrates may not be restricted to a single cortical locus, but distributed across several cortical loci [Sterzer et al., 2009]. Strong competitive interactions could thus potentially result from the multiplicity of mutually-inhibited populations encoding alternative percepts.

In contrast, the dynamics of sensory representations leading to perceptual reversals

is comparatively slow, suggesting the existence of a separate adaptation mechanism, and destabilizing the current state of alternative perceptual representations. This is consistent with the results of perturbation experiments which reveal that in between successive reversals, the dominant visual appearance gradually weakens, and becomes increasingly prone to experimenter-induced reversals [Wolfe, 1984, Nawrot and Blake, 1989, Petersik, 2002, Kang and Blake, 2010]. Presumably, adaptation may be a *joint* and *differential* process, in the sense that while support for the currently dominant appearance ‘adapts’, support for the alternative (and currently suppressed) appearance ‘recovers’. When a critical difference in activity is reached, a perceptual reversal is initiated.

In addition to reversal behaviour, the stationary statistics of dominance durations (i.e. for continuous presentation with constant input-levels) have been investigated extensively. In previous chapters, we only discussed modulations of the average dominance durations from the perspective of the scaling property, and did not consider their precise dependence on stimuli strengths. Under variations of input-levels, average dominance duration obey peculiar and somewhat counter-intuitive properties, on which we now further elaborate. These properties, also known as Levelt’s propositions, are a hallmark of binocular rivalry, and have been studied in great amount of detail since their initial discovery [Levelt, 1965]. In particular, they have been the subject to several reformulations, motivated by the inclusion of new experimental evidence [Brascamp et al., 2006, Klink et al., 2008, Kang, 2009, Brascamp et al., 2015]. In Levelt’s seminal study, which focused on binocular rivalry, only a restricted range of stimulus contrasts had been explored and his propositions had not yet been tested for other types of perceptual rivalry. Under their latest form, only two out of the four original propositions were retained as the ‘*core laws of binocular rivalry*’ (see [Brascamp et al., 2015] for an comprehensive review), because Levelt’s propositions I and III are included in the reformulation of proposition II.

Here, for the clarity of the argument, we chose to formulate these as follows:

- Levelt’s proposition I (L1): increasing the contrast to one eye increases the

average dominance for this eye's stimulus, and *also* reduces the dominance of the other eye's stimulus.

L1 thus corresponds to the *sign* of modulations of average dominance durations.

- Levelt's proposition II (L2): increasing the *difference* between stimulus strengths modulates the average dominance of the stronger percept more than the average dominance of the weaker percept.

L2 thus corresponds to the *amplitude* of modulations of average dominance durations.

- Levelt's proposition IV (L4): increasing the *average* stimulus strength while keeping the difference constant will reduce the average dominance durations for both percepts.

Note that these propositions are not exclusive to binocular rivalry, and are consistent with results obtained for other types of multistable displays. For other types of rivalry however, the parameter corresponding to the strength of the stimuli, i.e. the equivalent of the 'contrast' in binocular rivalry, may be different. In the case of kinetic depth effect, for instance, one can instead vary the luminance of leftward or rightward moving dots [Klink et al., 2008].

Current models of multistable perception have generally been designed and tested on the basis of the experimental observations presented above. In addition, our study includes the observations presented in Chapter 2 (scaling property and shape of reversal time densities, see **Fig. 2.3**).

5.2 Competition, adaptation and noise

Current models of multistable perception implement perceptual reversals using a balance between three main elements: competition, adaptation, and noise. Competition and perceptual exclusivity are obtained by mutual-inhibition between attractor states representing alternative percepts. The mechanism implementing perceptual reversals combines fast stochastic dynamics in a double-well attractor landscape with slow deterministic modulations of the energy landscape (adaptation). Such models can reproduce reversal behaviour with realistic distributions of dominance durations [Lehky, 1988, Laing and Chow, 2002, Moreno-Bote et al., 2007, Shpiro et al., 2007, 2009]. For restricted ranges of input-levels combinations, they can also provide a good, albeit partial, qualitative and quantitative agreement with Levelt's original propositions [Laing and Chow, 2002, Wilson, 2007, Moreno-Bote et al., 2007]. In particular, to reproduce L2, it is generally assumed that one contrast is kept fixed and at a high level, and to reproduce L4, contrast levels are assumed to be equal. However, more recent formulations of Levelt's propositions suggest that all combinations of contrast levels should be tested [Klink et al., 2008, Brascamp et al., 2015]. For extended range of contrast modulations, they have also predicted deviations from L4 which have not yet been, and may be difficult to verify experimentally [Shpiro et al., 2007, Curtu et al., 2008, Seely and Chow, 2011, Brascamp et al., 2015].

A notable shortcomings of current models is that they do not comply to the scaling property, which may present a more challenging constrain than the overall shape of reversal time densities, which are generically reproduced by mutual inhibition models with noise and adaptation [Kim et al., 2006]. Although several variants of these models have been introduced and rigorously analysed, here we illustrate this important limitation by proposing a simplified description of their dynamics. Perceptual reversals correspond to a noise-driven transition, and occur when the current state has been sufficiently destabilised by adaptation. This can be approximated as an escape process from a potential well of slowly changing depth, or equivalently, across a moving threshold. Formally, we assume that the instantaneous escape probabil-

ity reflects normally distributed noise with mean zero and variance σ^2 . We then compute the instantaneous probability that the noise exceeds the distance to the (time-varying) threshold $\theta(t)$:

$$p_{\text{esc}}(t) = \frac{\nu}{\sqrt{2\pi}} \int_{(x_{ss}-\theta(t))/\sigma}^{\infty} e^{-x'^2/2} dx'$$

where ν is an escape rate. Dividing time into discrete intervals Δt , we approximate the FPT density $f(t)$ as

$$\ln f(t) = \ln [p_{\text{esc}}(n\Delta t)] + \sum_{i=0}^{n-1} \ln [1 - p_{\text{esc}}(i\Delta t) \Delta t],$$

and evaluate the moments numerically (see **Fig. 5.1**).

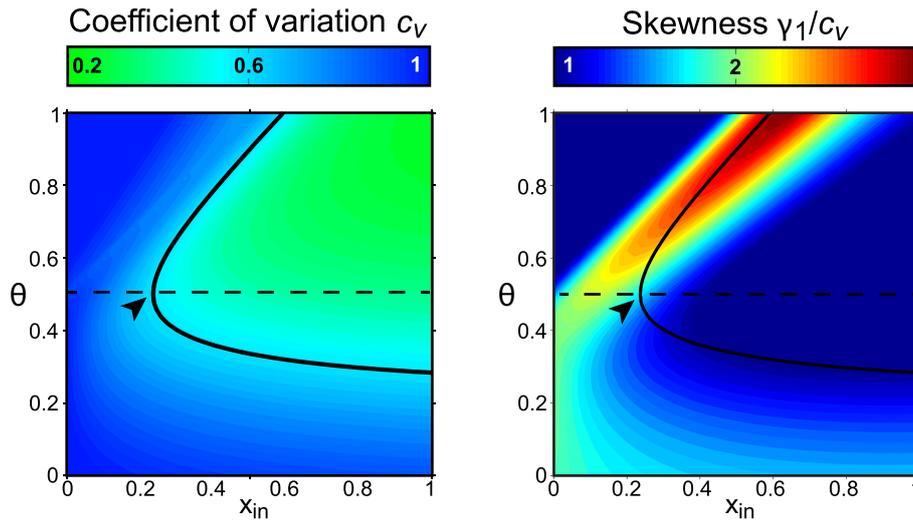


Figure 5.1: Adaptive threshold escape process: normalized moments of FPT density as functions of input x_{in} and threshold θ (both in units of σ). (*Left*): coefficient of variation c_v . (*Right*): skewness γ_1 (in units of c_v). Black curves mark $c_v = 0.6$. Black arrows indicate a particular combination of threshold and input-level for which a Gamma distribution is obtained. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

Fig. 5.1 shows that Gamma-like reversal times densities can be indeed be obtained by precisely matching adaptation rate, input strength, and noise amplitude, so that both the experimentally observed CV and skewness are reproduced (black arrow, $c_v \approx 0.6$ and $\gamma_1 \approx 2 c_v$). However, when input-levels are lowered, to obtain longer reversal times, adaptation no longer sufficiently assists perceptual

reversals and transitions become noise-dominated. FPTs become more and more irregular, causing the CV to rise significantly above the typical value observed in experiments. As a result, the shape of the distribution can not be maintained over the whole range of mean dominance durations imposed by experimental data, and the scaling property is not satisfied, as illustrated in **Fig. 5.2**.

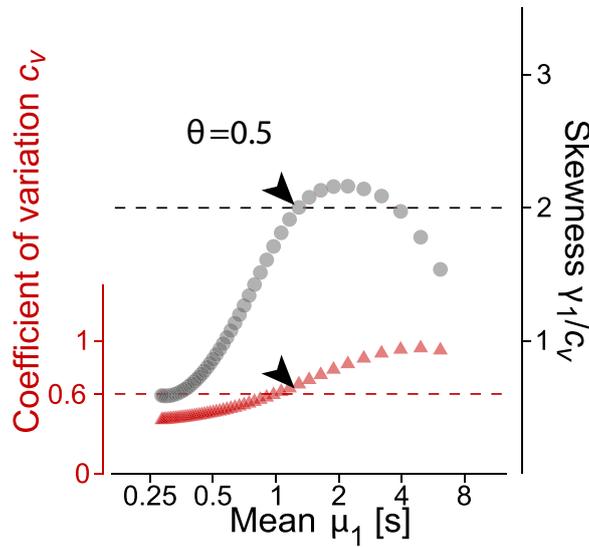


Figure 5.2: Adaptive threshold escape model: scaling property and FPT distribution shape as function of mean μ_1 (format as in **Fig. 2.3**). For constant θ (value given by inset), μ_1 decreases as x_{in} increases. The scalar property is satisfied for both the CV and the Skewness, but the ratio γ_1/c_v 3 is too large. Black arrows indicate a particular input-level for which a Gamma distribution is obtained. *Reprinted from [Cao et al., 2016] with the co-authors' permission.*

In what follows, we combine the knowledge gained from previous models with the new results established in previous chapters of this thesis, to propose a new model of multistable perception which simultaneously reproduces Levelt's proposition and both aspects of the scaling property (shape-invariance and Gamma-like shape of reversal time densities).

5.3 Qualitative hierarchical dynamics

To account for all order of the observed statistics of multistable perception, we propose to replace the slow, deterministic process of current models (adaptation) by a stochastic accumulation performed by the GE process. Indeed, we have established in Chapter 3 that the FPT statistics of the GE process could reproduce important properties of reversal time densities (scaling property and Gamma-like shape). We do not see a need to change the fast, stochastic processes of current models, which allow perceptual reversals to occur on a significantly shorter time-scale than typical reversal times.

This idea of a ‘noisy adaptation’ mechanism has also been suggested by several authors. For instance, an effect known as *stochastic resonance*: the time-course of perceptual reversals can be ‘enslaved’ to slow oscillations of input strength [Kim et al., 2006]. For an appropriate frequency (approximately half of the average dominance period), perceptual dynamics resonates and is driven near-deterministically by input oscillations. The same effect could be obtained in computational models available at the time of this last study (e.g. the stable oscillator and noise-driven attractor model [Lehky, 1988, Wilson, 2003, 2007]), with the important condition that some noise had to be introduced at the level of adaptation variables, not at the level of variables encoding perceptual states. A similar proposal has been made by van Ee, who has argued that reversal timing reflects a slow random walk rather than a deterministic relaxation, on the grounds that reversal timing is not entirely memoryless [van Ee, 2009].

The questions we now need to address are the following: how can we conciliate a slow dynamics of stochastic accumulation by the GE process, which reproduces the observed statistics of reversal time, with the fast dynamics implementing perceptual reversals? And how can we ensure that the interplay between both aspects of the dynamics will not compromise the peculiar dynamical properties on which the scaling property and the shape of reversal times densities rely?

A simple solution to these questions is to assume that neural populations perform-

ing the integration, or accumulation of sensory evidences, and those performing perceptual decision, are located at different levels in the hierarchy of visual processing. If accumulation and decision mechanisms are kept separate, and an appropriate structure of projections between levels is used, we can ensure that stochastic sensory accumulation leading to perceptual reversals will behave exactly as described in Chapter 3.

In the following section, we first present our new model of multistable perception from a qualitative angle. We describe the different levels of the hierarchy and illustrate their respective roles in perceptual dynamics. Next, we explain how the structure of interactions between the different constituents of the model is constrained by experimental observations, i.e. the scaling property, the Gamma-like shape of reversal time densities, and Levelt's proposition. Before introducing the model formally, from the perspective of its governing equations, we provide examples of its characteristic dynamics in situations relevant to our argumentation.

5.3.1 Levels of hierarchy

As previously mentioned, we assume two separate levels: an *evidence level* where sensory information is integrated, and a *decision level* where perceptual decision are performed. Both levels are populated by assemblies of bistable units, and therefore produce realizations of a GE process. At the evidence level, we assume bistable units to be independent - or weakly coupled - with individual transition rates directly mapped to stimulus strength. At the decision level, we assume strong recurrent connections between units as to obtain a global attractor property, i.e. so that the entire assembly can either be in a low or high activity state, with most units being inactive or active, respectively. To a given 'decision pool' (DP), we associate a perceptual representation, which we consider to be *dominant* when fully active, and *suppressed* when fully inactive. For each DP, we define, at the evidence level, an associated 'evidence pool' (EP), also composed of a finite number of bistable units, projecting excitatory, bottom-up projections to the DP. Global activation in a given DP, or equivalently, dominance of a given perceptual representation in awareness,

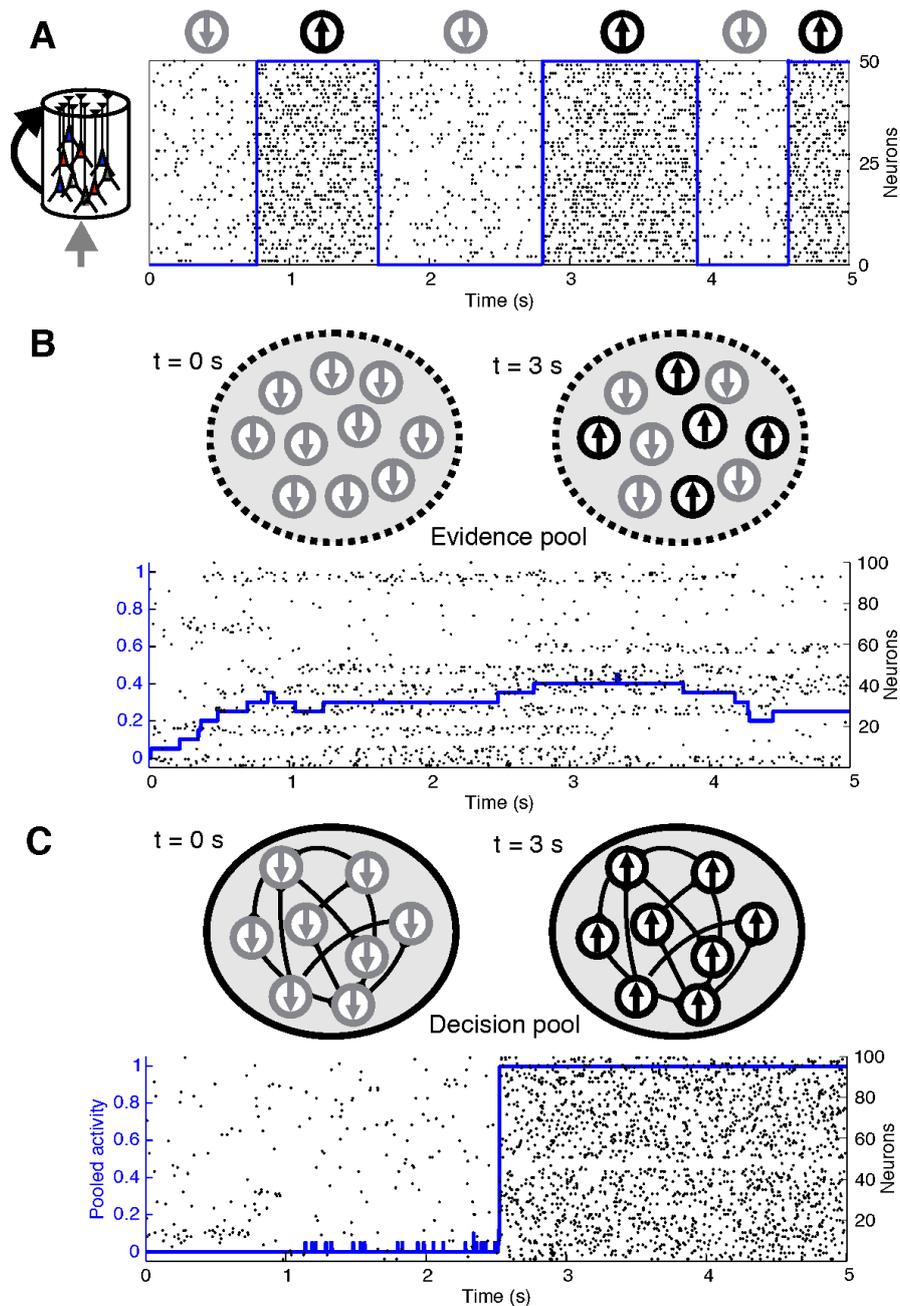


Figure 5.3: Model architecture (schematic). **A Microscopic scale** : A neuron assembly expressing ‘attractor dynamics’ (cartoon and spike raster) can be idealized as a *stochastically bistable unit* with activity (blue trace) transitioning spontaneously between ‘low’ and ‘high’ states (up and down arrows). **(bc) Hierarchy of representations:** Multiple bistable units can form *pools* to encode different levels of perceptual representations, from low-level stimulus components to high-level, structured perceptions. **B Mesoscopic scale Evidence pool (EP):** if units are unconnected, collective activity (blue trace) accumulates or dissipates stochastically in discrete steps, reflecting individual unit transitions. Spike rasters illustrate the activity of individual units (10 neurons per unit). **C Macroscopic scale Decision pool (DP):** if units are recurrently connected, collective activity (blue trace) is bistable with self-sustaining ‘low’ and ‘high’ states.

will depend on the number of active units in the associated EP. Additionally, we assume individual transitions rates in the DP to be large, compared to the transition rates in the associated EP. This results in slow accumulation of evidences to drive fast global activation in the associated DP, effectively implementing a hierarchical, FPT-like mechanism. These different levels of the hierarchy of perceptual representations used in our approach are illustrated in **Fig. 5.3** and correspond qualitatively to different scales of cognitive resolution. As such, they reflect three different scales of stochastic dynamics:

Individual bistable units, which idealize the dynamics of recurrently connected assemblies or neurons, or cortical columns, reflect sensory dynamics on a ‘*microscopic scale*’ (**Fig. 5.3 A**). Assemblies of bistable units within an EP, at the evidence level, reflect dynamics on a ‘*mesoscopic scale*’ (**Fig. 5.3 B**). There, accumulation of activity corresponds to the collectively evoked response of low-level, unstructured sensory representations of the stimulus’ features. Assemblies of bistable units within a DP, at the decision level (or behavioural level), reflect dynamics on a ‘*macroscopic scale*’ (**Fig. 5.3 C**). A DP possesses an all-or-none global attractor property, representing to the presence (when fully active) or absence (when fully inactive) in awareness of the associated representation of the stimulus. DPs can be thought of as encoding high-level, structured perceptual representations of the stimulus. When sufficient supporting evidences have been collected in an EP, the associated DP may transit to the active state (see joint dynamics in **Fig. 5.3 BC**).

5.3.2 Interactions within hierarchy

Here, we propose a brief overview of the structure of interactions in our model, which can be distinguished in two classes: competitive interaction between alternative representations, and hierarchical interactions between the evidence and decision level. Competitive interactions will appear at the decision level, and are nearly identical to that of other models, in that we use self-excitation and mutual inhibition to achieve categorical, mutually-exclusive perceptual decisions and fast reversals.

Hierarchical interactions feature both feedforward (bottom-up) and feedback

(top-down) influences. We consider two types of feedforward projections from the evidence level to the decision level: specific feedforward excitation, and non-specific feedforward inhibition. The former is implicitly illustrated in **Fig. 5.3**: specific feedforward excitation carries sensory evidences supporting a given representation (or pro-evidence), from a given EP to its associated DP, selectively, causing the corresponding representation to take dominance. Non-specific feedforward inhibition, on the other hand, carries non-supporting sensory evidences (or anti-evidence), from a given EP to all other competing DPs, non-selectively. In other words, a given percept can be positively biased by its corresponding set of sensory evidences, but can also be negatively biased by the set of sensory evidences corresponding to any other competing percept. The net flux between pro- and anti-evidence defines a relative evidence-bias, with important consequences for the resulting dynamics: perceptual choices and reversals will not solely depend on the absolute amount of evidence, but on the relative amount of evidence. Specifically, perceptual reversals will be initiated when a ‘differential threshold’ is reached, i.e. when the evidence supporting the suppressed representation sufficiently overcome that of the dominant representation.

We also consider specific feedback inhibition from the decision level to the evidence level, which is the key element implementing reversal behaviour. Each DP, upon taking dominance, selectively inhibits its associated EP. As a result, directly following a reversal, the EP supporting the newly dominant percept habituates, while the EP supporting the newly suppressed percept, and now free from such inhibition, recovers. This joint, differential and antiphase dynamics of sensory evidences will be central in several aspects of the discussion below.

Indeed, this implies that time intervals between successive durations depends on the time-course of both habituation and recovery processes, which themselves depend on contrast-levels. As a result, dominance durations will also directly depend on both contrasts, as required by Levelt’s propositions. Also, we can describe the time between reversals as the first-passage-time of a single variable, the evidence-bias (differential activity between EPs). Therefore, the important requirements underlying the scaling property discussed in Chapter 4, i.e. far-from-

equilibrium operation and low-threshold condition, can be easily transposed to our model's dynamics.

5.3.3 Qualitative description of dynamical behaviour

For the sake of clarity, we illustrate the characteristic dynamics of our model at the evidence and decision levels in the context of an unambiguous decision. We will then extend this circuitry to the case of ambiguous stimuli inducing multistable perception.

An unambiguous stimulus condition is modelled here by assuming that only one DP receives support from its associated EP at the evidence level, so that only one perceptual representation can be selected for this stimulus (**Fig. 5.4**).

When receiving sufficient support from its associated EP, a DP can become active, and the corresponding representation takes dominance in awareness. At the evidence level, other sets of bistable units, which do not necessarily support any specific higher-order representation, may also be activated by the stimulus and contribute to perceptual dynamics through non-specific feedforward inhibition. These projections are termed 'non-specific' in the sense that they are common to the whole decision level, not to a specific DP. This is to account for the fact that, for instance, noise or 'distractors' present in the stimulus may impair perceptual choices, by inhibiting the decision layer, thereby reducing its sensitivity. When evidence accumulation is somewhat conflicting, in the sense that it is not exclusively supporting a given representation, more evidences would have to be collected for a decision to be reached. **Fig. 5.4** also highlights the role of specific top-down inhibitory projections in implementing an habituation, or *perceptual adaptation* mechanism. Following a perceptual choice, the EP corresponding to the dominant percept are gradually suppressed. Note that this mechanism may not necessarily be engaged in truly unambiguous situations [Yuille and Kersten, 2006]. Here we exaggerate its effect for illustrative purposes, to highlight the similarity between this percept-driven habituation mechanism and the adaptation mechanism used in previous models, which will be of particular relevance in the case of an ambiguous

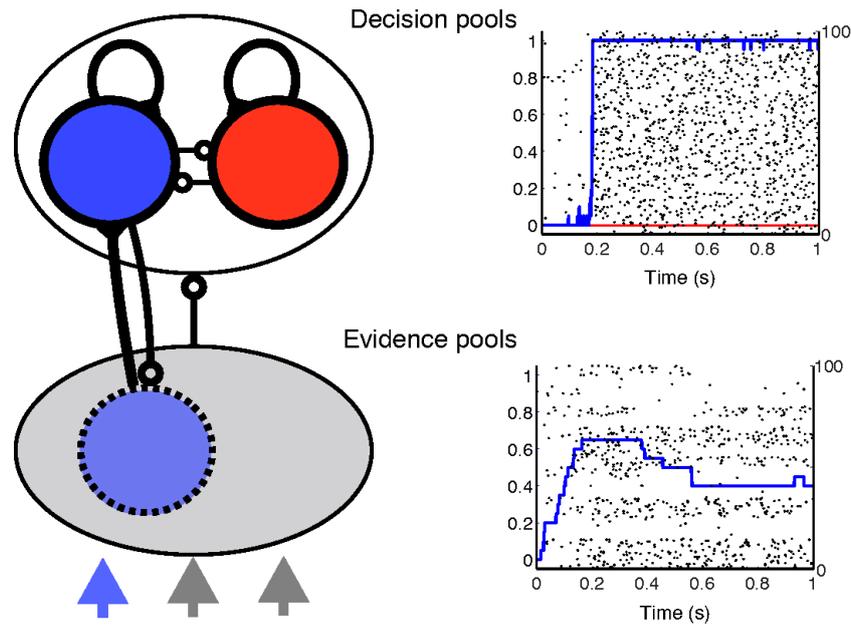


Figure 5.4: Perceptual dynamics at the macroscopic level: hierarchical interactions between evidence and decision pools reproduce perceptual dynamics of human observers. To account for perceptual adaption, EPs receive feedforward excitation from stimulus components (arrows) and feedback inhibition from the associated DPs. To reproduce an evidence-biased competition, DPs (dark blue, dark red) receive feedforward excitation from supporting evidence pools (light blue, light red), feedforward inhibition from other evidence units (gray), and horizontal inhibition from other decision pools.

Joint response to unambiguous stimulation: When stochastic accumulation of activity by evidence units (lower blue trace) reaches a differential threshold, this triggers global activation in the associated DP (upper blue trace) which, in turn, limits accumulation in the associated EP via feedback inhibition (lower blue trace). Spike rasters illustrate the activity of individual decision or evidence units (10 neurons per unit).

stimulus.

In the case where an unambiguous stimulus is presented, perceptual decision will - if at all - be unequivocal: accumulation of activity in the whole evidence level provides supporting evidences to only one possible interpretation of the stimulus. This simple detection circuit can be extended to the context of multi-stable perception and naturally produce percept alternations at the decision level. For this we only need to specify how the ambiguity of a stimulus translates in our model: some stimulus may drive stochastic accumulation not only in one, but in two distinct EPs, supporting two distinct and mutually-inhibited DPs.

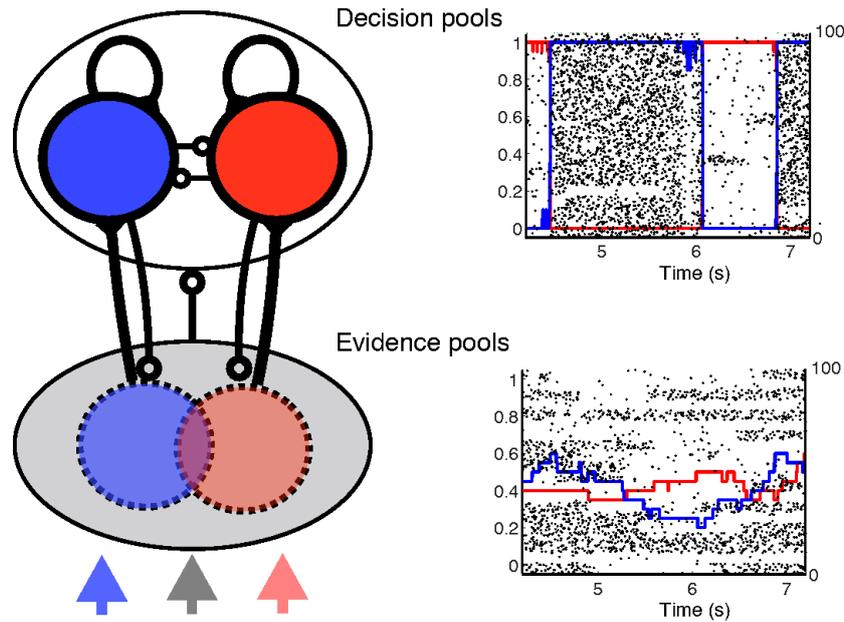


Figure 5.5: *Joint response to ambiguous stimulation* (details of hierarchical interactions and raster plots are as in **Fig. 5.4**): Competing stochastic accumulations of evidence units (lower blue and red traces) trigger simultaneous on- and offsets (perceptual reversals) in competing decision pools (upper blue and red traces). Each reversal shifts feedback inhibition from one evidence pool to another, reversing the previous accumulation trends in both pools (lower blue and red traces). The concomitant habituation/recovery phases are clearly visible between reversals. Each phase continues until differential evidence activity reaches a threshold.

Contrary to the unambiguous case, accumulation of activity at the evidence level is conflicting, in the sense that two mutually exclusive interpretations of the stimulus are simultaneously supported. This conflict between sensory evidences induces competition between the two percepts at the decision level. In this instance, the non-specific feedforward inhibition plays an important role. As previously mentioned, it implements an evidence-biased competition between representations. Perceptual dynamics is driven by differential accumulation of a ‘net flux’ of sensory evidence: reversals occur when a differential threshold of activity is reached at the evidence level (as opposed to an absolute threshold). In other words, competing appearances switch when the activity in one of the two EPs has increased sufficiently over the other (and vice-versa). The best supported decision pool then takes dominance, suppressing the other decision

pool via cross-inhibitory projections. Feedforward inhibition is commonly used in computational models of decision-making, in the context 2AFC tasks for instance [Mazurek et al., 2003, Bogacz et al., 2006, Larsen and Bogacz, 2010]. Some experimental evidence (local field potential recordings in monkeys) indicate that such ‘pro evidence’ (i.e. sensory evidences supporting a given decision) and ‘anti evidence’ (i.e. supporting a competing decision) are conveyed independently and that the net-flux is computed locally in higher-order association areas [Bollimunta and Ditterich, 2012].

Following a perceptual reversal, evidences supporting the newly dominant representation are forced to habituate downwards, driven by the dominant appearance through top-down inhibition. Simultaneously, evidences supporting currently suppressed representation can recover upwards: they are not subject to perceptual adaptation, simply because their associated DP is inactive. The differential activity threshold may be reached again, this time, in favour of the currently suppressed representation, and another perceptual reversal follows, starting the cycle anew (**Fig. 5.5**). It is important to note that while the initial percept selection is mostly driven the stimulus via bottom-up excitation, percept alternations are not solely stimulus-driven, but rather jointly driven by both bottom-up and top-down (percept-driven) signals. This supports the idea that multistable perception reflects hierarchical interactions between neural populations processing low- and high-level stimulus features [von Helmholtz, 1866, Leopold and Logothetis, 1999, Sterzer et al., 2009, Kang and Blake, 2010, Kornmeier and Bach, 2012].

In the proposed hierarchical structure, percept alternations naturally emerge when neural circuits implementing percept selection in unambiguous viewing conditions are subjected to ambiguous viewing conditions, where ambiguities in the stimulus can not be resolved. Here, conflicting stimulus integration at the evidence level drives perceptual dynamics at the decision level, and the statistics of reversal times directly reflect the statistics of stochastic accumulation in EPs. The present study distances itself from previous models in that adaptation is not implemented as an external (‘ad-hoc’) and passive mechanism, but rather as

an internal and active mechanism, emerging from the structure of interactions within the model's hierarchy. In adaptation-based models, adaptation variables are usually adiabatically coupled to decision variables: when one takes dominance, the associated adaptation current slowly builds up, gradually weakening the dominant representation, until a perceptual reversal is initiated by external noise. Importantly, the strength of adaptation currents evolve on a fixed time-scale, are set and tuned by an external parameter, and only indirectly depend on input-levels. A direct consequence is that longer time-scales obtained experimentally (i.e. longer than the time-scale of adaptation itself) can not be reached without violating the scalar property, because the system invariably moves from an adaptation-driven regime (comparable to the DDR we have described in Chapters 2-3), to a NDR.

Our mechanism of perceptual adaptation, implemented by a joint dynamics of evidence accumulation and of evidence habituation, resolves these issues. It is equivalent to an adaptation in that it drives perceptual alternations by reversing the net drive received by each decision pool after a new percept has taken dominance, weakening the dominant representation and allowing the suppressed representation to eventually take dominance again. The EP supporting the newly dominant percept habituates downwards, driven by top-down, percept-specific feedback inhibition, therefore reducing the amount of supporting evidences. Meanwhile, the EP supporting the newly suppressed percept recovers upwards, driven by bottom-up feedforward excitation, therefore increasing the amount of supporting evidences. The difference is that here, the strength of habituation and accumulation 'currents' implementing perceptual adaptation do not only depend on fixed projections weights, but also directly depend on input-levels: both recovery and habituation evolve towards a steady-state which varies with stimulus intensity (see **Fig. 5.6**). Furthermore, through the individual transition rates of bistable units, the characteristic time-scale of EPs also depends on input-levels. This allows for dominance durations to be mutually dependent on both stimulus strength, and will constitute the basis for the explanation of Levelt's propositions in a future section. To be more precise, the recovery of suppressed sensory evidences is *faster*

when the corresponding contrast *increases*, so that the dominance time of the currently dominant percept will be shortened. Conversely, the habituation of dominant sensory evidences is *slower* when the corresponding contrast *increases*, so that the dominance time of the currently dominant percept will be lengthened.

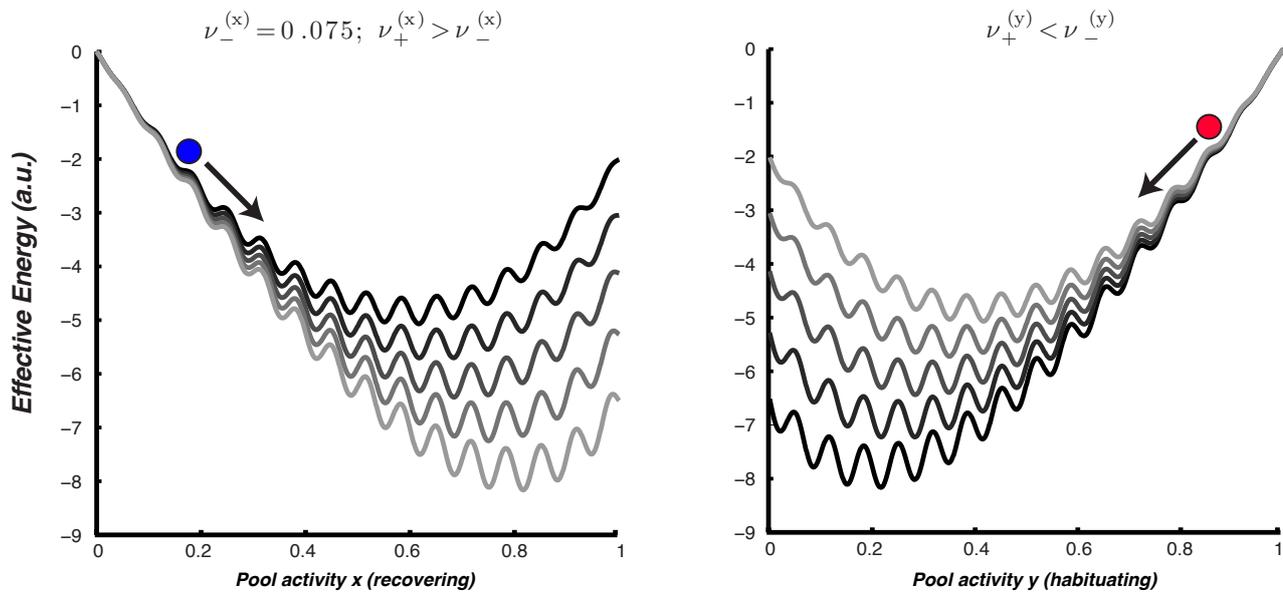


Figure 5.6: Schematic illustration (arbitrary units) of the joint dynamics of sensory evidences in their respective energy landscapes, or ‘nested attractor’, following a reversal where percept X (blue) as been suppressed and percept Y (red) has taken dominance. When a sufficient difference in activity is reached (differential threshold), another reversal occurs, and the situation is reversed. (*Left*): Activity in evidence population x (blue), supporting percept X . Lighter hues indicate larger contrast value c_x (*Right*): Activity in evidence population y (red), supporting percept Y . Lighter hues indicate larger contrast value c_y .

To summarize, in our model, perceptual reversals occur as a natural by-product of mutual interactions between sensory representations, and categorical representations. The dynamics of sensory evidences drives perceptual decisions in a bottom-up manner, and, in a top-down manner, current perceptual states cause sensory evidences to reorganise. Directly following a reversal, the decision fully agrees with the evidence, in the sense that the ‘best’ supported percept - with the most associated evidence - is indeed dominant, and the other, suppressed. Due to specific feedback inhibition, the dominant percept then gradually loses its

support, while the suppressed percept regains its support. This eventually results in an increasingly compelling contradiction, where the current perceptual state no longer agrees with the evidence, which culminates in the form of a perceptual reversal.

5.4 Formal hierarchical dynamics

5.4.1 Dynamical equations

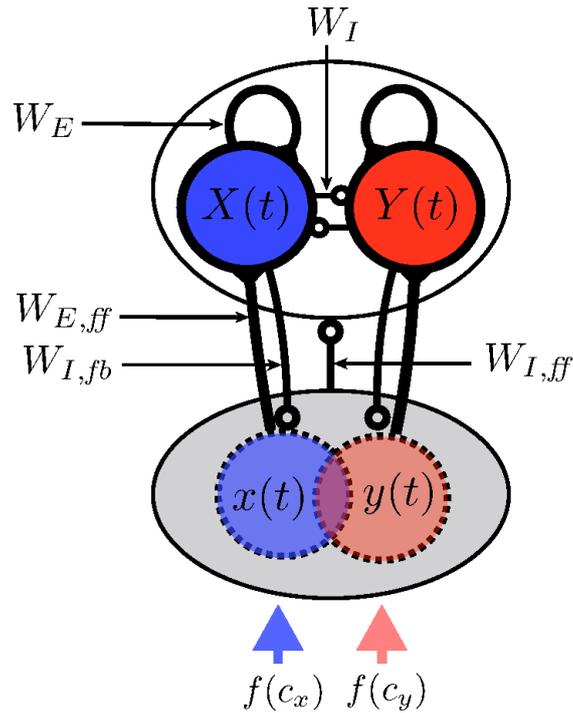


Figure 5.7: Structure of the hierarchical model

In this section, we detail the equations governing our model's dynamics. All populations in our model (EPs or DPs) obey the Master Equation for the GE process (3.2). For a pool indexed by "p" of N_p stochastic nodes with rates $\nu_{\pm}^{(p)}$, the temporal evolution of the probability $P_{n_p}(t)$ to find n_p active nodes at a given time

t is:

$$\begin{aligned} \frac{dP_{n_p}(t)}{dt} = & \nu_+^{(p)} [N_p - (n_p - 1)] P_{n_p-1}(t) + \nu_-^{(p)} (n_p + 1) P_{n_p+1}(t) \\ & - [\nu_+^{(p)} (N_p - n_p) + \nu_-^{(p)} n_p] P_{n_p}(t) \end{aligned} \quad (5.1)$$

In all simulations, we set the number of units in each pool to 25. For individual switching rates, we assume that spontaneous transitions of bistable units correspond to an escape process (as depicted in **Fig. 3.1**), so that they obey an Arrhenius-like equation (as in Kramer's escape problem):

$$\nu_{\pm}^{(p)} [U_p(t)] = \frac{1}{2\tau_p} e^{\pm \frac{U_p(t)}{2}} \quad (5.2)$$

Where we define $U_p(t)$ the total incoming activity to pool "p", which varies overtime, and will depend on the network's interaction constants. For example, if $U_p(t) = 0$, both energy wells in **Fig. 3.1** would have the same depths, and $\nu_+ = \nu_-$. If $U_p(t) < 0$, because of the opposite sign in expression 5.2, $\nu_+ < \nu_-$ and the energy well corresponding to the inactive state is deeper than the one corresponding to the active state, because. Therefore there will be more downward transitions than upward transitions (same as **Fig. 3.1**). Also, the characteristic time-constant for each pool is defined for as:

$$\tau_p^{(0)} = \left[\nu_+^{(p)}(0) + \nu_-^{(p)}(0) \right]^{-1} \quad (5.3)$$

The instantaneous time-constant at time t for each pool is therefore:

$$\begin{aligned} \tau_p(t) &= \left[\nu_+^{(p)}(U_p(t)) + \nu_-^{(p)}(U_p(t)) \right]^{-1} \\ &= \tau_p^{(0)} \operatorname{sech} [U_p(t)] \end{aligned} \quad (5.4)$$

And the instantaneous asymptotic state x_{in} takes the form of a sigmoid function:

$$\begin{aligned}
 x_{in} &= \frac{\nu_+^{(p)}(U_p(t))}{\nu_+^{(p)}(U_p(t)) + \nu_-^{(p)}(U_p(t))} \\
 &= \frac{1}{1 + \exp\left(-\frac{U_p(t)}{2}\right)} \\
 &= \phi(U_p)
 \end{aligned} \tag{5.5}$$

Where we have defined a scaled sigmoid function $\phi(x) = (1 + \exp(-2x))^{-1}$

We now define population activities for the evidence level with lower-case letters: $x(t) = \frac{n_x(t)}{N_x}$ and $y(t) = \frac{n_y(t)}{N_y}$. Associated decision populations are identified with upper-case letters: $X(t) = \frac{n_X(t)}{N_X}$ and $Y(t) = \frac{n_Y(t)}{N_Y}$.

Incoming activity to pools (x, y, X, Y) are defined as at any time t by the following relations:

$$U_x(t) = f(c_x) - W_{I,fb}X(t)$$

$$U_y(t) = f(c_y) - W_{I,fb}Y(t)$$

$$U_X(t) = W_{E,dec}X(t) - W_{I,dec}Y(t) + W_{E,ff}x(t) - W_{I,ff}(x(t) + y(t)) + \theta_{dec}$$

$$U_Y(t) = W_{E,dec}Y(t) - W_{I,dec}X(t) + W_{E,ff}y(t) - W_{I,ff}(x(t) + y(t)) + \theta_{dec}$$

This corresponds to a mean-field approximation, where we consider that the incoming activity to each pool of bistable units is the weighted sum of afferent average activities, with the value of weights correspond to the synaptic coupling between interacting pools (see **Fig. 5.7**). In total, the model consists of 11 independent parameters, which we summarise here:

- τ_{int} : EPs characteristic time-constant.
- τ_{dec} : DPs characteristic time-constant.
- $W_{E,dec}$: excitatory self-coupling of decision pools, which replaces recurrent connectivity between bistable units in DPs.

- $W_{I,dec}$: mutual inhibition between decision pools, corresponding to the fact competing representations are mutually exclusive.
- $W_{E,ff}$: specific bottom-up excitation received by decision pools, from their associated supporting evidence pools.
- $W_{I,fb}$: specific top-down inhibition received by supporting evidence pools, from their associated decision pools.
- $W_{I,ff}$: non-specific bottom-up inhibition from the evidence level to the decision level.
- θ_{dec} : an offset parameter to the decision layer, setting the shape of the energy landscape for zero-input.

For instance, negative values of θ_{dec} corresponds to a deeper well for the inactive state, so that the decision pool remains inactive when no stimulus is present.

- $f(c_{x,y})$: evidence layer input, where c_x and c_y correspond to the ‘contrast’, or stimulus strength to evidence pools x and y respectively.

For the optimisation procedure to behavioural data described in a later section, a logarithmic input-mapping was used: $f(c) = \alpha \log(c + \gamma) + \beta$, with c , the contrast.

5.4.2 Analysis of decision level: perceptual threshold and switching behaviour

To understand how the perceptual decision threshold is implemented at the decision level, and can be reached by accumulation of activity in the evidence layer, we need to look at the how the fixed points of a given decision pool evolve under such accumulation. For this, we can use the deterministic approximation of the Master

Equation for one of the decision pools, for instance X :

$$\begin{aligned}\frac{dX(t)}{dt} &= \nu_+^{(X)} - \left(\nu_+^{(X)} + \nu_-^{(X)}\right) X(t) \\ \tau_X(t) \frac{dX(t)}{dt} &= \phi [W_E X(t) - W_I Y(t) + W_{E,ff} x(t) - W_{I,ff} (x(t) + y(t)) + \theta_{dec}] - X(t) \\ &= \phi [W_E (X(t) - U_X(t))] - X(t)\end{aligned}$$

Where we have set, $U_X(t) = \frac{1}{W_E} (W_I Y(t) - W_{E,ff} x(t) + W_{I,ff} (x(t) - y(t)) + \theta_{dec})$. Therefore, $X(t)$ will be drawn towards one of the fixed point(s), X_{ss} defined as the solutions of the following equation:

$$X_{ss} = \phi [W_E (X_{ss} - U_X(t))] \quad (5.6)$$

Depending on the value of $U_X(t)$, and assuming $W_E > 2$ (so that the slope of function ϕ in zero is greater than one), this equation can have one or two stable fixed points: a low activity fixed point where $X_{ss} \approx 0$, and a high activity fixed point where $X_{ss} \approx 1$. $U_X(t)$ corresponds to the position of the inflexion of the *transfer function* ϕ .

The competing decision pool Y satisfies the same equations where variables X and Y as well as x and y have been exchanged.

Starting from the low-activity fixed point, we obtain the reversal threshold by determining the value of the evidence-bias which would trigger a perceptual reversal, i.e. the value $\Delta(t) = x(t) - y(t)$ for which the low-activity fixed point is destabilised, so that the activity of the decision pool $X(t)$ must switch to the high-activity state. This situation is illustrated in **Fig. 5.8**.

The value of $U_X(t)$ for which the low-activity fixed point disappears, U_\uparrow can be obtained numerically, for a given value of the self-excitatory weight W_E . Defining the perceptual threshold θ as the value of $\Delta(t)$ so that $U_X(t) = U_\uparrow$, and $Z(t) = (x(t) + y(t))/2$, the average activation at the evidence level, we obtain:

$$\theta = \frac{2}{W_{E,ff}} (W_{I,dec} - \theta_{dec} - U_\uparrow W_{E,dec}) - \frac{2}{W_{E,ff}} (W_{E,ff} - 2W_{I,ff}) Z(t) \quad (5.7)$$

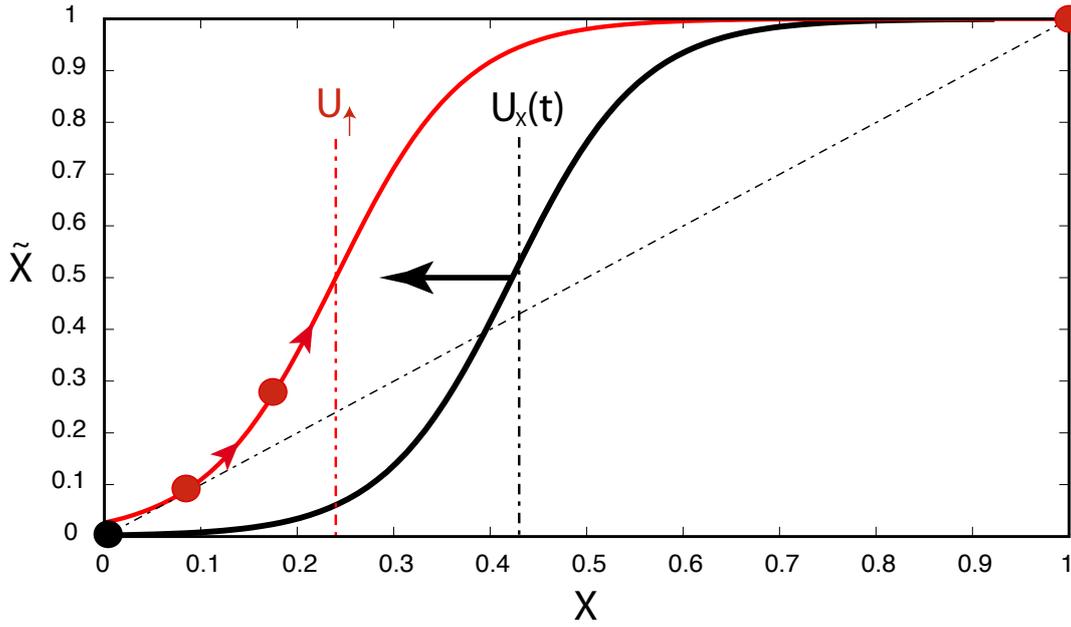


Figure 5.8: Destabilization of the low-activity fixed point: Dotted diagonal line corresponds to the identity function $\tilde{X} = X$. Thick black line corresponds to the function $\tilde{X} = \phi[X - U_X(t)]$. Thick red line corresponds to the function $\tilde{X} = \phi[X - U_\dagger]$. Thick left-point arrow corresponds to the shift of the activation function, driven by the accumulation of activity in the evidence pool, i.e. when the evidence bias $\Delta(t) = x(t) - y(t)$ increases. For $U(t) > U_\dagger$, the activation in pool X is near zero (black dot). For $U_X(t) < U_\dagger$, the low-activity fixed point vanishes, and the decision pool activity is drawn towards the high-activity fixed point (arrowed red dots).

The perceptual threshold thus has a constant component and a time-dependent component, so that $\theta = \theta_0 - \theta_1 Z(t)$ with:

$$\theta_0 = \frac{2}{W_{E,ff}} (W_{I,dec} - \theta_{dec} - U_\dagger W_{E,dec}) \quad (5.8)$$

$$\theta_1 = \frac{2}{W_{E,ff}} (W_{E,ff} - 2W_{I,ff}) \quad (5.9)$$

Note that expression (5.7) does not take into account the fluctuating dynamics in DPs, and thus overestimates the value of effective reversal thresholds. This is because as $U(t) \rightarrow U_\dagger$, and the low-activity attractor becomes shallower, finite-size fluctuations may drive the decision pool activation over the energy barrier which separates it from the high-activity attractor; as a result, the effective threshold will be smaller than expected in our deterministic approximation.

However, expression (5.7) is useful in that it exposes the dependence of the threshold on the average activity $Z(t)$ in the evidence pools, which will be of relevance in later sections of this chapter. In particular, for the set of parameters which provides the optimal fit to the experimental data, we obtain $W_{E,ff} > 2W_{I,ff}$, and thus $\theta_1 > 0$, so that the threshold for perceptual reversals tends to decrease with $Z(t)$. This introduces a degree of adaptive sensitivity at the decision level: if the average level of evidence increases, the reversal threshold decreases. The decision layer thus becomes more sensitive as activity increases, because a smaller evidence-bias is required for a perceptual reversal to ensue.

5.4.3 Analysis of evidence level: evidence-bias

One important feature of our model of multistable perception is that the joint dynamics of accumulation and habituation, analogous to adaptation currents in other models is input-dependent. This is because the increment and decrement rates in either the recovering or habituating pools depend on the contrast strengths c_x and c_y . In this section, we propose a deterministic toy-model for the FPT of the evidence-bias $\Delta \doteq x(t) - y(t)$, which will prove useful in understanding the properties of average dominance durations under stimulus modulations (and in particular, Levelt's proposition II). For the sake of clarity, let us assume that population X is currently inactive (the associated representation is therefore suppressed) and population Y is currently active (the associated representation is therefore dominant), so that the corresponding evidence pools x and y are recovering upwards, and habituating downwards respectively.

Specifically, the evidence pool supporting the currently suppressed percept recovers towards an asymptotic value x_{in} and the evidence pool supporting the currently dominant percept habituates towards an asymptotic value y_{in} . x_{in} depends on the

contrast level $c_x = c_{sup}$ and y_{in} depends on the contrast level $c_y = c_{dom}$

$$x_{in}(c_{sup}) = \frac{\nu_+^{(x)}(c_{sup})}{\nu_+^{(x)}(c_{sup}) + \nu_-^{(x)}(c_{sup})} \quad (5.10)$$

$$y_{in}(c_{dom}) = \frac{\nu_+^{(y)}(c_{dom})}{\nu_+^{(y)}(c_{dom}) + \nu_-^{(y)}(c_{dom})} \quad (5.11)$$

In the deterministic limit, the evolution equations for x and y are simple exponential relaxations towards their respective asymptotic values, given by (5.10) and (5.11), with their respective time constant τ_x and τ_y :

$$\tau_x \frac{dx}{dt} = x_{in} - x \quad (5.12)$$

$$\tau_y \frac{dy}{dt} = y_{in} - y \quad (5.13)$$

We further reduce this description by assuming that the habituation and recovery processes exactly mirror each other, i.e. that $\nu_+^{(x)} = \nu_-^{(y)} \doteq \nu_+$ and $\nu_+^{(y)} = \nu_-^{(x)} \doteq \nu_-$. Using (5.12) and (5.13), we obtain the following equation for the dynamics of the evidence-bias:

$$\tau_\Delta \frac{d\Delta(t)}{dt} = \Delta_{in} - \Delta(t) \quad (5.14)$$

$$\tau_\Delta = \frac{1}{\nu_+ + \nu_-} \quad (5.15)$$

$$\Delta_{in} = \frac{\nu_+ - \nu_-}{\nu_+ + \nu_-} = x_{in}(c_{sup}) - y_{in}(c_{dom}) \quad (5.16)$$

Where we have introduced the quantity Δ_{in} , the asymptotic evidence-bias.

Perceptual reversals occur when the evidence bias reaches a threshold θ . Recall that here, we have assumed X to be suppressed and Y to be dominant, and defined the evidence bias accordingly: during the dominance of Y , $\Delta(t)$ evolves from $-\theta$ to θ (during the dominance of X , $\Delta_{in} \rightarrow -\Delta_{in}$ and $\Delta(t)$ evolves from θ to $-\theta$). Therefore, assuming $\Delta_{in} > \theta$, the average FPT for the evidence-bias, or equivalently, the duration of the current dominance, T_{dom} , is simply that of an exponential relaxation:

$$T_{dom} = \tau_\Delta \ln \left(\frac{\Delta_{in} + \theta}{\Delta_{in} - \theta} \right) \quad (5.17)$$

The duration of the current dominance phase T_{dom} obeys a simple relation to stimulus strength, with three important properties.

Firstly, it depends on *both* contrasts, through the asymptotic evidence bias Δ_{in} and the characteristic time τ . This is because the accumulation of activity driving reversals is *differential*, therefore the current dominance depends as much on the recovery time of the EP supporting the suppressed percept as on the habituation time of the EP supporting the dominant percept (see caption in **Fig. 5.9 AB**).

Secondly, it depends *non-linearly* on Δ_{in} , meaning that for a given change of the contrasts, the corresponding modulations of T_{dom} may have different amplitudes depending on the value of Δ_{in} is small (i.e. when T_{dom} is long, and the corresponding percept is *strong*), or large (i.e. when T_{dom} is short and the corresponding percept is *weak*). Thirdly, the value of the differential reversal threshold θ also depends on both contrasts, because the average activity at the evidence level $Z(t)$ does (see also, (5.7)).

Importantly, Δ_{in} , mostly depends on the *difference* between contrasts, while θ mostly depends on the *sum* of contrasts. This observation will be useful when discussing the mechanism of Levelt's propositions, because L2 corresponds to relative variations of contrasts (and thus to modulations of Δ_{in}) and L4 corresponds to absolute variations of contrasts (and thus to modulations of θ).

Naturally, the assumption required to obtain (5.17), namely that habituation and recovery processes exactly mirror each other is generally inadequate. However, as will be shown in the next sections, this expression is sufficient to qualitatively capture important aspects of Levelt's propositions.

5.5 Agreement with behavioural data

5.5.1 Model optimization

Here, we provide some details on the optimization procedure carried out to fit the model's statistics to the available behavioural data. The data initially consisted of the mean dominance times μ_1 , the CV c_v and the ratio between the skewness and

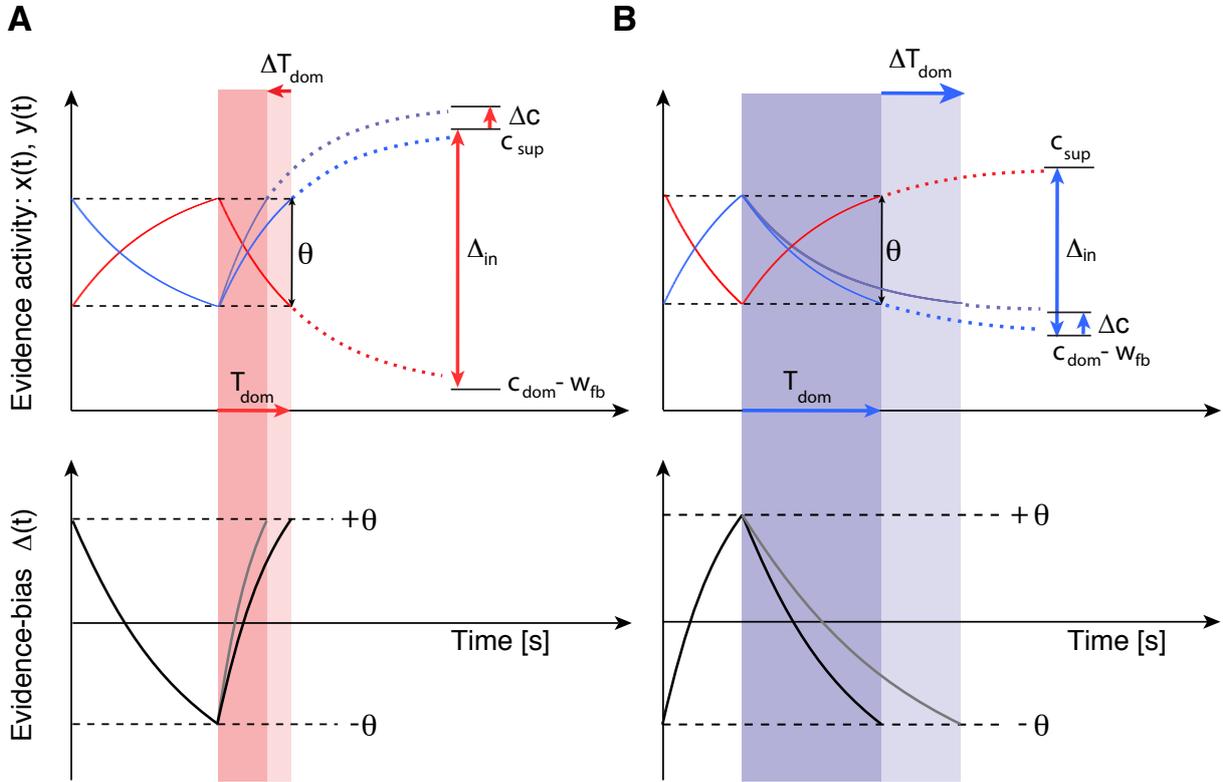


Figure 5.9: AB Deterministic evidence activities (top, red and blue traces) and differential activity, or evidence-bias (bottom, black traces), over two successive dominance periods. Activities recover, or habituate, exponentially until a differential threshold θ is reached. Dotted curves extrapolate the exponential time-courses to their respective steady-states, which reflect input, c_{sup} (recovering evidence), or input minus feedback inhibition, $c_{dom} - w_{fb}$ (habituating evidence). Dominance durations are controlled by the range between steady-states, Δ_{in} : **A** When the contrast of the perceptually suppressed input (here, corresponding to blue traces) is incremented by Δc , Δ_{in} increases and T_{dom} decreases. **B** When the perceptually dominant input is incremented by a same amount Δc , Δ_{in} decreases so that T_{dom} increases.

CV γ_1/c_v , obtained for a binocular rivalry display. Note that due to the noisiness of experimental data in terms of the skewness, and to ensure the convergence of the fit, we later assumed that the relation $\gamma_1 = 2c_v$ was exactly satisfied. For each observer, 25 values were collected for each quantities, corresponding to a 5×5 array of contrast combinations between the following values: [6.25%, 12.5%, 25%, 50%, 100%]. Average dominance times were individually normalized, and averaged over observers. Higher-order moments were averaged over observers. In addition, the average correlation between successive durations (at lag 1) over the whole dataset was included in the optimisation procedure, for the following reason: although generally modest, the correlation reported in experiments is consistently positive (as reviewed in [van Ee, 2009]), and alternative statistical measures have been introduced which establish a subtle but robust history-dependence in alternation series ('cummulative history' [Pastukhov and Braun, 2011]). This new line of evidence will be further confirmed in a later section of this chapter, we will further contribute to this new line of evidence with theoretical predictions corroborated by new observations.

The experimental dataset therefore consisted of three 5×5 arrays ($X_i^{(exp.)}$, $i \in [1, 3]$ for the mean, CV and skewness of the average observer), and a scalar $\bar{X}_4^{(exp.)}$ for the correlation, with \bar{X} the matrix average of X . The corresponding simulated dataset $X_i^{(sim.)}$, $i \in [1, 3]$ (for the moments) and $\bar{X}_4^{(sim.)}$ (for the correlation) were constructed accordingly. Alternation series were generated from the model by averaging over 10 trials of 2 minutes durations each, to match experimental conditions. The fit error, or mean residual \bar{R} , was computed as follows:

$$\bar{R} = \frac{1}{4} \sum_{i=1}^4 w_i \bar{\delta}_i$$

$$\delta_i = \left| \frac{X_i^{(exp.)} - X_i^{(sim.)}}{\bar{X}_i^{(exp.)}} \right| \quad i \in [1, 3]$$

$$\delta_4 = \left| \frac{\bar{X}_4^{(exp.)} - \bar{X}_4^{(sim.)}}{\bar{X}_4^{(exp.)}} \right|$$

A total of around 400 minimizations of the mean residual \bar{R} starting from random initial configurations of the model parameters were performed using a

stochastic gradient descent algorithm [Spall, 2003]. We used the following weighting $w = [1, 1, 1, 1/4]$, to prioritize the fit of moments surfaces and ensure some correlations were obtained for the optimised parameter set of the model.

Note that we did not fit the entire ‘correlation-surface’. This is an important detail, because although we required the model to produce some correlations, we did not specify the structure of correlations observed in the data.

The average fit error \bar{E} is related to the average residual \bar{R} by:

$$\bar{E} = \frac{\sum_i w_i \bar{\delta}_i}{\sum_i w_i} = \frac{16}{13} \bar{R}$$

For the optimal parameter set, this resulted in an average fit error, of about 15% with the ‘mean observer’ dataset. Supplementary details on the optimization procedure, including the optimal set of parameters, can be found in Appendix B.

5.5.2 Fit results: average dominance durations

Fig. 5.10 illustrates the agreement between experimental data, and simulations of our model, using an optimized parameter set (for details on parameters optimization, see B). The quantitative match between simulations and experiments is remarkable for all combinations of contrasts, illustrating that our model complies to all of Levelt’s propositions.

5.5.3 Levelt’s Proposition I and II

Our optimized model not only provides quantitative account of L1 and L2 (**Fig. 5.10 A**), but also allows a straightforward qualitative insight in their origins. Since the model always operates in a DDR, alternations are mainly driven by deterministic factors and we can study the properties of average reversal times by neglecting fluctuations, i.e. we assume that the number of bistable units is extremely large. In this case, the accumulation of activity at the evidence level reduces to the toy-model we described in 5.4.3, and use expression (5.17) to characterise the average duration of the current dominance.

L2 states that increasing stimulus strength to one eye modulates the dominance

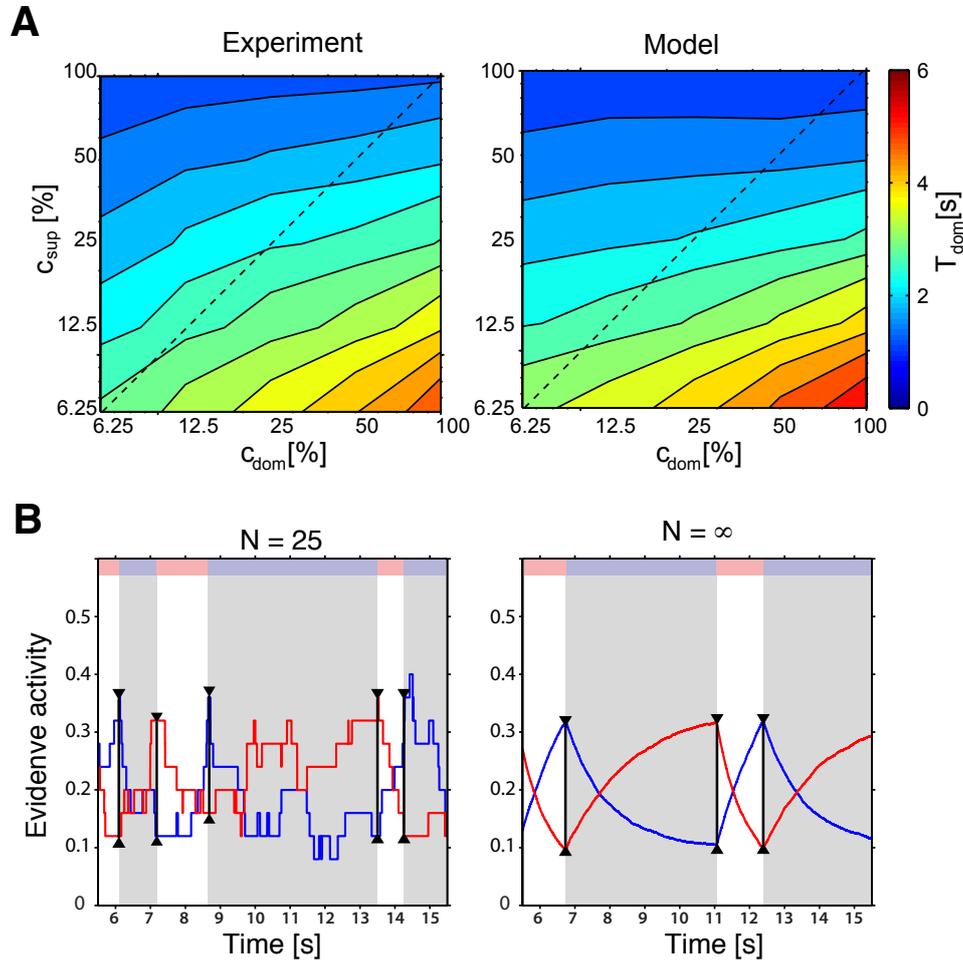


Figure 5.10: Mean dominance durations vary characteristically with dominant and suppressed stimulation (“Levelt’s propositions”, L1, L2, L4). Our model explains these observations by concomitant habituation of dominant evidence, and recovery of suppressed evidence, until a differential threshold is reached. **A** Mean dominance duration T_{dom} increases with contrast c_{dom} of dominant input and decreases with contrast c_{sup} of suppressed input, in both experiment (*left*) and model (*right*). Both trends are exacerbated when suppressed input is weaker ($c_{sup} < c_{dom}$). Dotted lines mark the equidominance condition ($c_{sup} = c_{dom}$) **B** Activity of evidence pools (red and blue traces) over several dominance periods. Stochastic reversals with variable thresholds θ (black lines) of finite system, with $N = 25$ units per pool (*left*) and deterministic reversals with fixed θ of infinite system, with $N \rightarrow \infty$ (*right*).

duration of *both* percepts, and also that if one of the two percepts is predominant, modulations will be of greater amplitude for this percept than for the weaker one. We therefore consider the situation in which the relative perceptual dominance is unbalanced, so that one percept can be considered *strong* (with a long average dominance duration) and the other one, *weak* (with a short average dominance duration) (**Fig. 5.10 B**).

The first aspect of Levelt's propositions is that changing the contrast to one eye changes the average dominance for both percepts, in opposite directions (L1). This feature naturally follows from the architecture of the model, because perceptual switches are driven by the difference of activity between EPs. Recall that, following a reversal, the EP associated with the newly dominant percept habituate downwards, while the EP associated with the newly suppressed percept recovers upwards. The next reversal is initiated when the differential threshold is reached again. As a result, a given dominance duration depends as much on the time-course of the habituation phase in its supporting EP as on the recovery phase in the conflicting EP (i.e. the EP supporting the alternative percept). From expression (5.17), we see that average dominance durations, assuming a fixed reversal threshold θ , mainly depends on the value of the asymptotic evidence bias Δ_{in} , which itself depends on both input-levels c_{sup} and c_{dom} : it increases with c_{sup} and decreases with c_{dom} . This means that increasing the input-level to the EP supporting the currently suppressed percept accelerates its recovery and therefore shortens the current dominance T_{dom} (see **Fig. 5.9 A**). Conversely, increasing the input-level to the EP supporting the currently dominant percept slows down its habituation and therefore lengthen the current dominance T_{dom} (see **Fig. 5.9 B**), consistently with L1.

The second aspect of Levelt's propositions is that modulations of average dominance durations have different amplitude depending on whether the dominant percept is strong or weak. This is well captured by the approximated expression of T_{dom} 5.7 obtained in the deterministic reduction of our model. At the evidence level, dur-

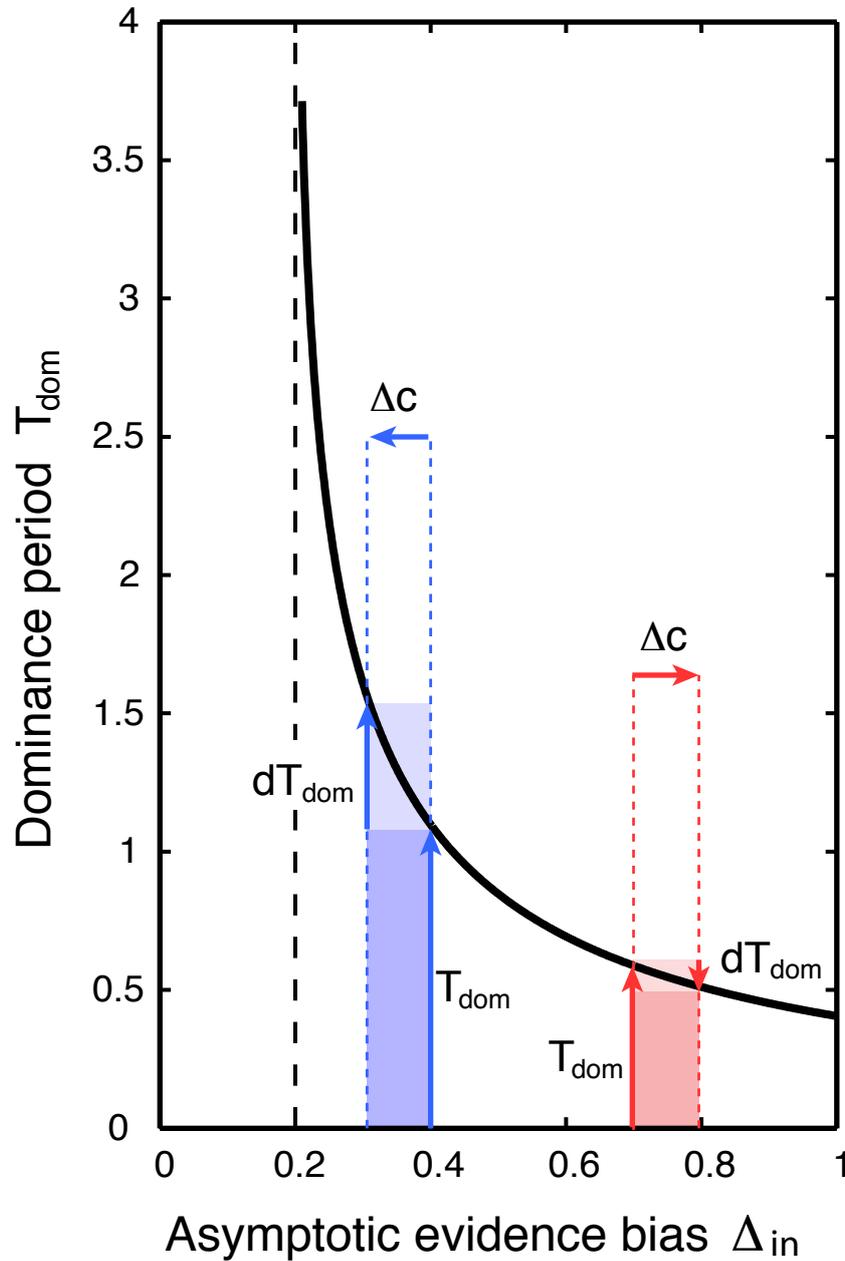


Figure 5.11: Modulations of average dominance T_{dom} with the asymptotic evidence bias Δ_{in} explain L2. Enlarging the range Δ_{in} accelerates the growth of differential activity and curtails dominance periods T_{dom} (and *vice versa*), see Equation (5.17). Because the dependence is hyperbolic, differential changes to *smaller* ranges Δ_{in} entail *larger* effects on T_{dom} (and *vice versa*).

ing the dominance of a stronger percept, both recovery and habituation processes are slow, because the asymptotic evidence-bias Δ_{in} is small, and therefore, T_{dom} is long. Conversely, during the dominance of a weaker percept, both recovery and habituation processes are fast, because the asymptotic evidence-bias Δ_{in} is large, and therefore, T_{dom} is short. Also, the dependence of T_{dom} on Δ_{in} is hyperbolic: as a result, stronger percepts show larger modulations than weaker ones, under the same change in input strength, consistently with L2 (see Eq. 5.7, **Fig. 5.9 AB** and **Fig. 5.11**). Other models have reached similar qualitative conclusions: they usually assume adaptation to follow an exponential profile, which yields the same asymmetry in modulations of dominance duration between weak and strong percepts [Brascamp et al., 2015]. Approximated expressions for the average dominance duration have previously been obtained by several authors [Laing and Chow, 2002, Wilson, 2007]. These expressions are formally similar to the expression (5.17) we obtained in Section 5.4.3, in that they all reflect the first-passage-time of an exponentially relaxing variable. Here, it is the evidence-bias $\Delta(t)$ which relaxes exponentially towards Δ_{in} (**Fig. 5.9 AB**, bottom).

Quantitatively, our model provides a complete account for Levelt's propositions: it was fitted to experimental data for all combinations contrasts. In other models, L2 is obtained by varying only one of the two contrasts, while the other one remains fixed (at high level), providing a good albeit partial fit [Laing and Chow, 2002, Wilson, 2007].

Note that, although we have not explicitly considered the input-dependence of the characteristic time-constants for both recovery and habituation processes it does not conflict with the overall picture described here. Increasing the input-level for c_{sup} and c_{dom} respectively shortens and lengthens the time-scale of the corresponding processes, and thus affect average dominance durations in the same direction than variations off Δ_{in} , only to a lesser extent. This can also be seen from Eq. (5.17): modulations of T_{dom} are linear in τ and hyperbolic in Δ_{in} .

5.5.4 Levelt's Proposition IV

As illustrated by **Fig. 5.10 A**, not only L1 and L2, but also L4 are well reproduced by our model. The qualitative explanation of L4, however, differs from that proposed by other models [Moreno-Bote et al., 2007, Wilson, 2007, Shpiro et al., 2007, Curtu et al., 2008, Seely and Chow, 2011]. In adaptation-based models, variations of inputs induce modulations of dominance durations by directly acting on the relative stability of percepts. For instance, increasing contrasts together tends to make the suppressed percepts less stable, yielding greater reversal rates. However, to some extent, it also tends to make the dominant percept more stable. A direct consequence of this is that modulations of dominance durations under variations of input strength are not always strictly monotonic and may deviate from L4 at low contrast levels (i.e. reversal rates decrease for increasing contrasts, rather than increase). This is because high- and low-contrast reversals may be prompted in two different manners: they can be initiated by the upward switching of the suppressed percept ('escape'), or by the downward switching of the dominant percept ('release') Curtu et al. [2008]. Compliance of these models to L4 generally require extra-assumptions to balance the system's sensitivity between adaptation currents and input-levels, such as a restriction or pre-processing of inputs (as thoroughly reviewed in [Seely and Chow, 2011]).

These results led to the proposal that L4 may in fact no longer be valid for low-contrasts and called for additional experimental data. However there is some evidence of that L4 is valid down to detection levels [van Ee, 2009]. It has also been noted that such deviations from L4 may be challenging to verify, as they would typically occur in a narrow range close to detection thresholds, making it difficult to sample experimentally [Brascamp et al., 2015].

The data presented in this study, in both the experiment and the model, also seem to support the validity of L4 (see **Fig. 5.10**, left). Here, we propose an alternative mechanism to account for these results. An important feature of the model is that modulations of average dominance durations do not reflect a *static*, but a *dynamic* change of the energy landscape. In previous models, changes in

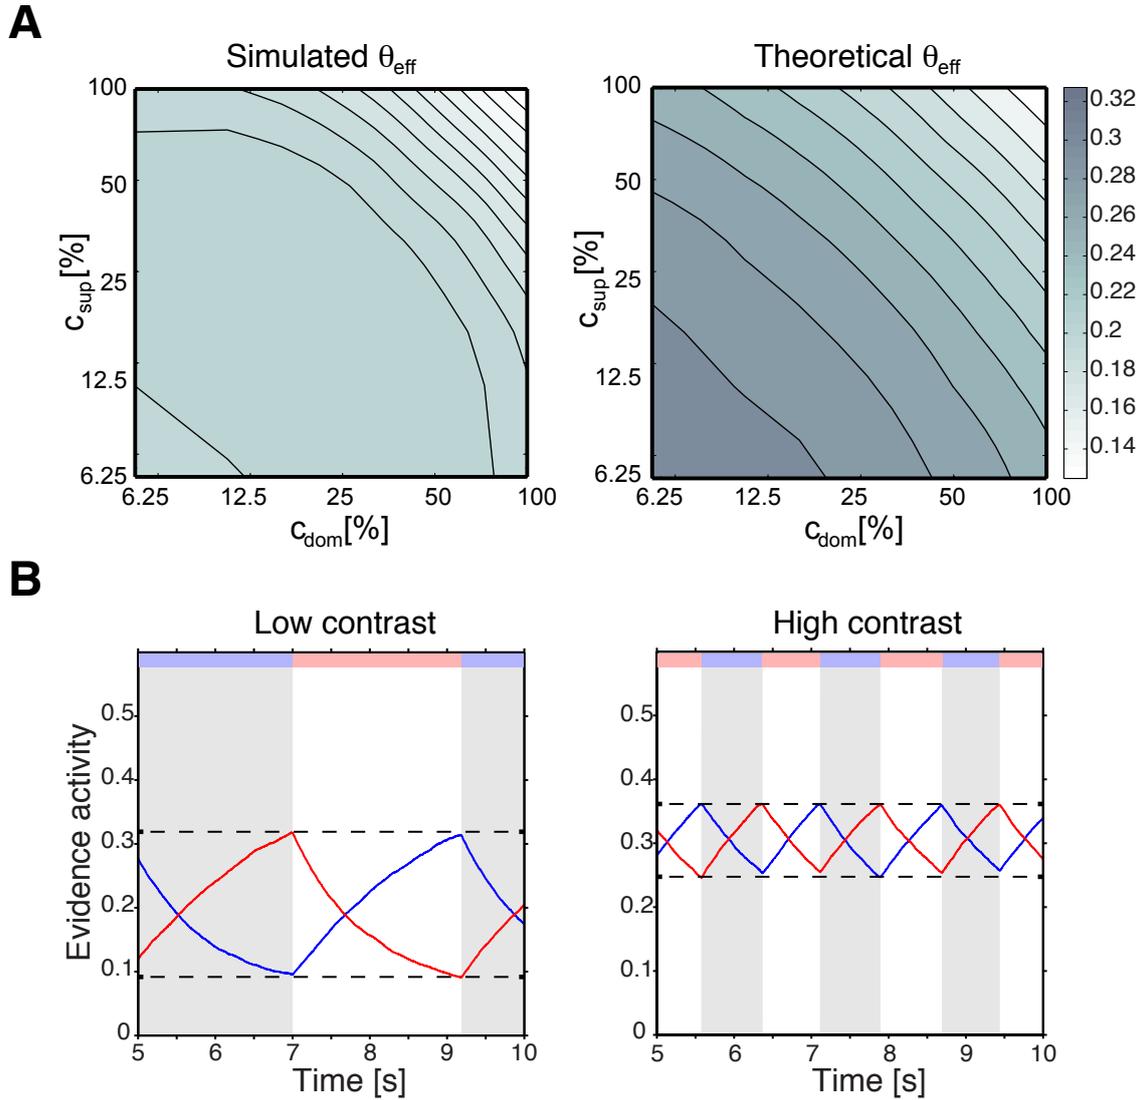


Figure 5.12: Mean dominance durations decrease with stimulation (“Levelt’s 4th proposition”). Our model explains this observation in terms heightened competition between decision pools reducing differential thresholds for evidence pools. **A** Average differential evidence activity in model simulations, just prior to reversals, as a function of input contrasts c_{dom} and c_{sup} . **B** Theoretical reversal threshold θ_{eff} of decision populations (Eq. 5.7), as a function of input contrasts. **C** Average activity of evidence pools (solid traces) over several dominance periods (deterministic system with $N \rightarrow \infty$). Left: for low input contrast, evidence activity recovers and habituates over a *lower* and *broader* range (dashed lines), increasing dominance periods. Right: for high input contrast, evidence activity traverses a *higher* and *narrower* range (dashed lines), decreasing dominance periods.

input-levels directly affect the stability of dominant and suppressed appearances by shaping the depth of the energy wells (steady-states), making them more or less stable. In contrast, in our model, the stability of decision pools between reversals does not change. Rather, changes in input-levels affect the time-course of the habituation and recovery processes, which results in modulations of average dominance periods and explains L2. In the case of L4, we suggest that a different mechanism may be at work.

We have mentioned in Section 5.4.3 that the differential threshold has a constant and a time-varying component, $\theta = \theta_0 - \theta_1 Z(t)$. This means that variations of the average stimulus strength modulate the sensitivity of the decision level: depending on the level of activity at the evidence level (which itself depends on the input-levels), the value of the reversal threshold changes. From the approximate expression of the perceptual threshold obtained from the deterministic limit of the model (see (5.7)), we have learned that the effective value of the perceptual threshold depends on the average activity $Z(t)$ at the evidence level. This dependence is weighted by a pre-factor $\theta_1 = W_{E,ff} - 2W_{I,ff}$, which corresponds to the net balance between feedforward excitation (provided by each evidence pool to its associated decision pool), and feedforward inhibition (provided by the whole evidence level to the whole decision level). In particular, for our optimal parameter set, $\theta_1 > 0$: as $Z(t)$ varies overtime, the effective evidence-bias required for a perceptual reversal to ensue, will also vary overtime. This is an important aspect of the model's dynamics which will be further detailed in a later section of this chapter, when discussing the origins of the small but consistent sequential correlation between successive dominance durations.

Also, when the sum of either stimulus strength increases, the average activity $\langle Z \rangle$ increases. This is simply because when doing so, asymptotic values for both evidence pools increase, shifting the average level of activity upwards. Therefore, the average threshold $\langle \theta \rangle$ (e.g. over an entire trial), decreases with increasing stimulus strength. Shorter reversal thresholds naturally yield shorter dominances, consistently with L4.

This means the decision layer becomes more sensitive as the amount of total evidence increases because a lesser evidence-bias is needed to trigger a perceptual switch (**Fig. 5.12**).

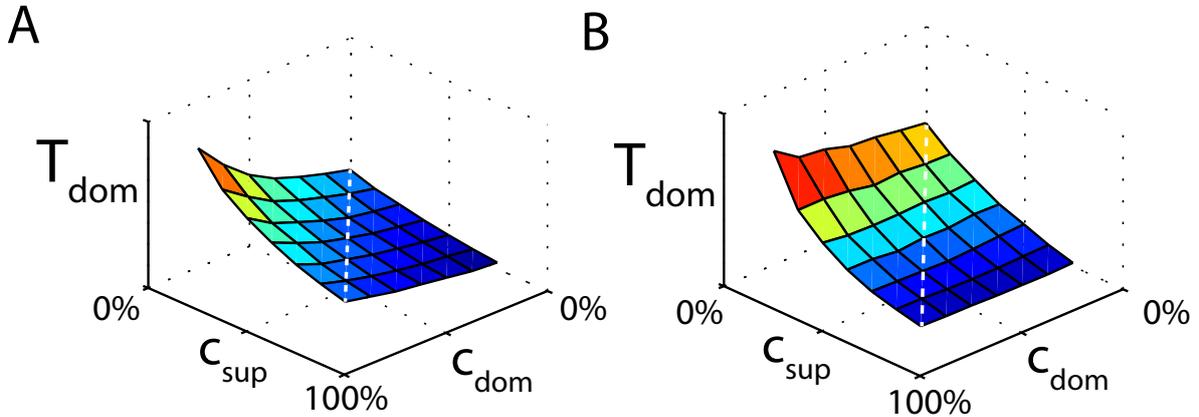


Figure 5.13: Dominance time modulations in the context of Levelt’s fourth proposition arise from threshold modulations. Dotted white lines mark the equidominance $c_{dom} = c_{sup}$. **A** $W_{E,ff} = 2W_{I,ff}$ so that the threshold no longer depends on the average activity Z in the evidence pools. As a result, we observe no modulations of dominance durations along the equidominance. **B** $W_{E,ff} > 2W_{I,ff}$ so that the threshold varies with the average activity Z . We observe visible modulations of the average dominance duration along the equidominance.

To further confirm the role of threshold modulations in reproducing L4, we can examine the changes of the average threshold $\langle \theta \rangle$ with the average evidence activity $\langle Z \rangle$, in two different cases: the case where $W_{E,ff} = 2W_{I,ff}$, and the case $W_{E,ff} > 2W_{I,ff}$. In the former case, $\theta_1 = 0$, and as a result, we do not observe modulations of dominance durations for a simultaneous increase of both contrasts **Fig. 5.13 A**. In the latter case, $\theta_1 > 0$ and a decrease of average dominance durations as both contrasts increase are clearly **Fig. 5.13 B**.

Note that although the characteristic relaxation times for both recovery and habituation processes do depend on input-levels, time-scale modulations do not significantly contribute to L4. This is because, as previously mentioned, an increase of both contrast c_{dom} and c_{sup} have opposite effects on the time-scale of their respective processes, and essentially cancel each other out: suppressed evidences have a faster

recovery, but this is balanced by dominant evidences having a slower habituation. Therefore, our model suggests that L4 may correspond to changes in the sensitivity of observers to perceptual reversals.

Note that situations where L4 is not observed (as in **Fig. 5.13 A**), and where L4 is particularly salient (as in **Fig. 5.13 B**), have both been reported experimentally (see [Moreno-Bote et al., 2010] and [Kang, 2009], respectively). Interestingly, it has been suggested that both situations may reflect differences in how incoming stimuli strengths are normalized [Brascamp et al., 2015]. Our results support this recent conjecture in that a continuous transition between these two cases is obtained by modulating the relative contribution of specific feedforward excitation (relaying the activity of *one* EP) and of non-specific feedforward inhibition (relaying the activity of the *sum* of the two EPs).

5.5.5 Fit results: scaling property and distribution shape

In this chapter we have extended the reduced one-dimensional FPT framework developed in Chapter 2-4, to two dimensions, where two competing percepts explicitly compete for dominance. We have previously established that the GE process, when operating in a drift-dominated regime, could satisfy all aspects of the scaling property [Cao et al., 2016]. This requires, in particular, that the threshold remains small compared to its asymptotic value, to ensure accumulation proceeds sufficiently far from equilibrium. The structure of interactions within the model carefully preserves this regime, so that for most combinations of input-levels our model satisfies the scaling property, along with the characteristic Gamma-like shape.

This is because the process driving perceptual reversals, the evidence-bias, is simply the difference of two antiphasic GE processes, so that the time between successive reversals corresponds to its FPT to the perceptual threshold θ . Accumulation of evidence remains in a drift-dominated regime for all combinations of contrasts, and reversal thresholds remain small compared to the asymptotic evidence-bias ($\theta \ll \Delta_{in}$). Even in the case of extremely unbalanced contrasts, where changes in the CV indicate the proximity of the noise-dominated regime,

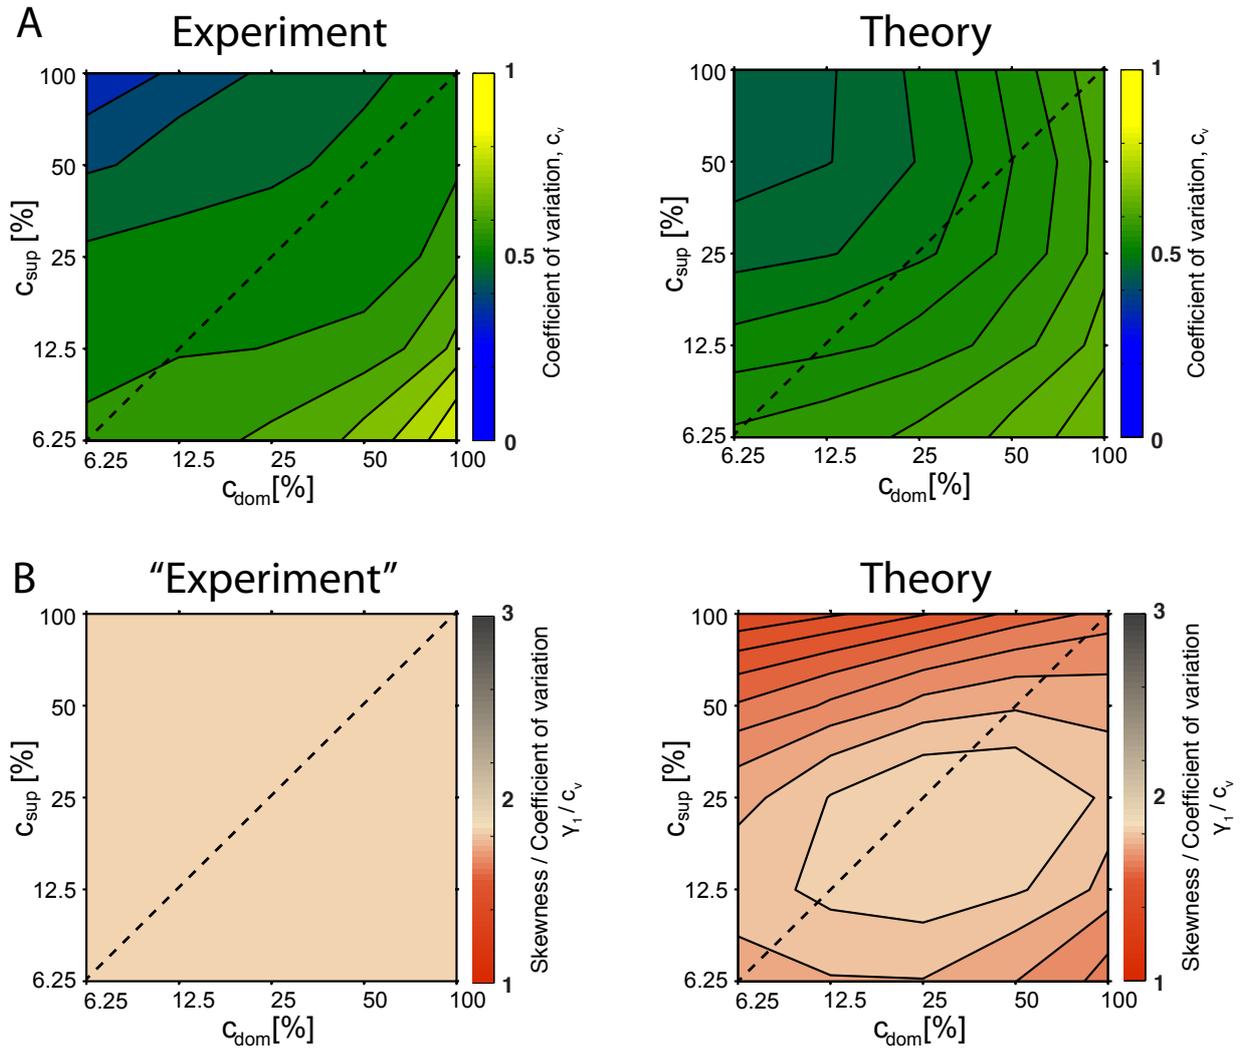


Figure 5.14: The distribution shape of dominance periods remains almost unchanged over a wide range of input contrast (‘scaling property’), implying that input proportionally impacts drift and noise of an underlying stochastic accumulation. **A** Coefficient of variation of dominance times T_{dom} as a function of input contrasts c_{dom} and c_{sup} , in both experiment (*left*) and model (*right*). For symmetric input contrasts (dashed line), the coefficient of variation remains essentially unchanged at $c_v \approx 0.5$. Variability changes only for highly asymmetric contrasts, decreasing when $c_{sup} \gg c_{dom}$ and increasing when $c_{sup} \ll c_{dom}$. **B** Ratio between the skewness and CV of dominance times T_{dom} as a function of input contrasts c_{dom} and c_{sup} . Note that due to the noisiness of experimental data in terms of the skewness, and to ensure the convergence of the fit, we have assumed that the relation $\gamma_1 = 2c_v$ was exactly satisfied. As for the CV, the skewness in simulated alternation series remained essentially unchanged, for symmetric input contrasts (dashed line), at $\gamma_1 \approx 2c_v$. For highly asymmetric contrasts, $c_{sup} \gg c_{dom}$ or $c_{sup} \ll c_{dom}$, the skewness slightly decreased.

threshold-crossings at the evidence level remain consistently sharp. The sharpness of reversals at both levels of the hierarchy clearly appears when observing the average behaviour of activity before and after reversals (see transition-triggered averages in Appendix B). As a result, our optimised model fits well the statistics of behavioural data at second-order (CV) and third-order (skewness). The scaling property is obtained, and reversal time densities are close to a Gamma distribution, as illustrated in **Fig. 5.14**).

We now conclude this section by making some additional remarks.

Note that the scaling property of multistable perception is not strictly satisfied for all contrast combinations. In particular, in both experiments and simulations, the CV changes most for relative modulation of contrasts (i.e. changes in the difference between contrasts): weak percepts have greater CV than average, and strong percepts have lower CV than average (see **Fig. 5.14 A**). As previously mentioned, this could suggest vicinity of the boundary with a single fixed point regime, where only one percept is reported and reversals no longer occur. On the other hand, we there are no significant variations of the CV when both contrasts increase simultaneously, in seemingly apparent contradiction with the results established in Section 5.5.4. Indeed, if this causes the threshold to gradually shrink, we should also expect the CV to increase, as would be the case in the one-dimensional FPT problem discussed in previous chapters.

However, in this case we consider the FPT of $\Delta(t) = x(t) - y(t)$, which consists two independent GE processes, for which the variance is the sum of the variances of process $x(t)$ and $y(t)$. This effectively increases the value of the threshold required to obtain $c_v \approx 0.6$, so that small variations of the threshold do not affect the CV as severely as for a one-dimensional process. However, the analysis of the FPT problem for $\Delta(t)$, i.e. obtain the FPT moments in the (Δ_{in}, θ) -space, would be required to confirm this intuition.

The scaling property is also not strictly satisfied for the ratio of skewness and CV, in similar conditions, i.e. when contrasts are most unbalanced (see **Fig. 5.14 B**). In these conditions, the resulting reversal time densities becomes closer to a Weibull

distribution (with $\gamma_1 \lesssim 2 c_v$). As noticeable in our initial observations (see **Fig. 2.3 C**), experimental data for the ratio of skewness and CV is considerably more distributed around its average. In particular, for the binocular rivalry experiment, it spreads widely across the average value, within the range $c_v \lesssim \gamma_1 \lesssim 3 c_v$. In the case of the moving plaids, a fraction of the data lies well below the average value, with $c_v \lesssim \gamma_1 \lesssim 2 c_v$. While further data collection would be required to draw definite conclusions at this order of the statistics, it may be interesting to verify whether the data for which $\gamma_1 \simeq c_v$ corresponds to strongly unbalanced situations, as predicted by our model.

5.6 Non-stationary dynamics

We have shown in previous sections that discrete stochastic dynamics, embedded in a hierarchical model, can reproduce reversal behaviour, as well as all known characteristics of the stationary statistics of alternation series obtained experimentally. Quantitative agreement with Levelt's propositions, the scaling property and the stereotypical shape of reversal time densities is remarkable, and simple qualitative mechanisms which may explain these important empirical observations have been proposed.

Even in stationary conditions, i.e. continuous viewing and constant stimulus qualities within trials, the model predicts interesting non-stationary behaviour. In particular, slow modulations of the stochastic dynamics at the sensory level introduce some degree of short-term memory at the perceptual level. In this section, we propose a preliminary study of non-stationary properties in alternation series produced by the model and confront our predictions to additional analysis of behavioural data.

5.6.1 Correlations and structure of correlations

Until recently, the general consensus was that successive dominance and suppression periods were statistically independent and sequences of perceptual alternations were 'memoryless', and sequential correlations recorded in experiments were deemed too

small to be significant [Fox and Herrmann, 1967, Borsellino et al., 1972, Walker, 1975, Lehky, 1995, Pastukhov and Braun, 2007]. Recent studies have put this into question, by reporting consistently small but significant correlations [van Ee, 2009, Pastukhov and Braun, 2011, Pastukhov et al., 2013]. Adaptation-based models generally produce correlated alternation sequences, except when adaptation is not sufficient to drive perceptual reversals. In this case reversals are mostly noise-driven, and correlations can be sufficiently reduced. This led to the proposition that adaptation does not trigger perceptual reversals alone, but destabilises the steady-state enough to allow noise-driven transitions to occur [Brascamp et al., 2006, Moreno-Bote et al., 2007]. However, we have shown that setting such balance between adaptation and noise would eventually conflict with the scaling property, for instance when attempting to obtain longer average dominance durations. Moderate correlations can also be obtained when dominance durations are modulated by slow changes of a ‘memory’ component of the dynamics, as proposed in [Gigante et al., 2009].

Here, we propose that correlations may be explained by an alternative mechanism.

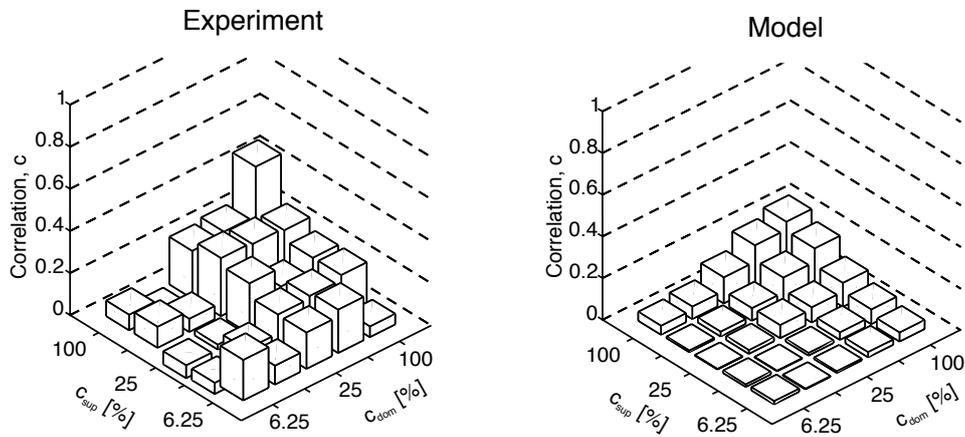


Figure 5.15: Successive dominance periods become correlated at sufficiently high stimulus contrast. Correlation coefficient c of successive dominance periods increases with contrast, c_{dom} and c_{sup} , in experimental (*left*) and simulated data (*right*).

As illustrated in **Fig. 5.15**, the model predicts, in some cases, a small but significant correlation coefficient between successive dominance durations (i.e. lag 1 correlation coefficient). For both observed and simulated series of dominance du-

rations, the serial correlation coefficient was computed as follows: we consider a sequence $\{T_i\}$ of dominance durations, and extract the sequence $\{T_i^{(\alpha)}\}$ of dominances corresponding to percept $\alpha \in [X, Y]$, i.e. one or the other percept, as well as the sequence $\{T_i^{(\alpha^*)}\}$ of the directly following dominances. We also note $T^{(\alpha)}$ the average dominance of the sequence $\{T_i^{(\alpha)}\}$ and $T^{(\alpha^*)}$ the average dominance of the sequence $\{T_i^{(\alpha^*)}\}$. The correlation c is then obtained as:

$$c = \frac{\langle (T_i^{(\alpha)} - T^{(\alpha)})(T_i^{(\alpha^*)} - T^{(\alpha^*)}) \rangle}{\sqrt{\langle (T_i^{(\alpha)} - T^{(\alpha)})^2 \rangle \langle (T_i^{(\alpha^*)} - T^{(\alpha^*)})^2 \rangle}} \quad (5.18)$$

As explained in Section 5.5, the correlation obtained in experimental alternation sequences was also included in the optimization of the model's parameters, therefore the existence of overall correlations should not be surprising. However, the model was only constrained to the average correlation over the whole range of contrast not to the whole 'correlation-surface', unlike the moments (mean, CV and skewness), for which a 5×5 array was used. A smaller weight was used for the correlation than for moments, to prioritize the fit of moments' surfaces. This means that although the optimised model was tuned to produce some correlations, we made no extra assumption with respect to the structure of these correlations, and how they may depend on contrasts variations. In particular, the model predicts such correlation to be near zero at low contrast, and to increase with contrast, reaching a maximum near 0.2 in the maximum contrast condition. This trend was corroborated by additional observations on the available experimental data. That correlations would only become apparent only in high contrast conditions, when dominance times are short, may be the sign of an intrinsic short-term memory, with a characteristic time-scale shorter than average dominance durations obtained for moderate and low contrast values.

We next examined how successive dominance periods within an alternation sequence were paired together. If successive dominance durations were statistically independent, and their correlation equal to zero, then the sequence of dominance durations T_i should be well described by a renewal process. Formally, this means that the joint distribution between a given dominance duration T_i and the next, T_{i+1} , should

factorize as follows:

$$P(T_i, T_{i+1}) = P(T_i)P(T_{i+1}) \quad (5.19)$$

If successive dominance durations, however, carry some degree of statistical dependence, the joint probability will not factorize. To assess whether the sequence dominance durations generated by our model, and those observed in experiments correspond, or depart from the statistical of a renewal process, we define the ‘correlation map’ ΔP as follows [Schwalger, 2013]:

$$\Delta P = P(T_i, T_{i+1}) - P(T_i)P(T_{i+1}) \quad (5.20)$$

ΔP should be identically zero if a sequence of dominance durations was generated by a renewal processes, and if not, it should reveal the underlying structure of correlations, which add up to the correlation coefficient. Another relevant quantity is the conditional expectation $\langle T_{i+1}|T_i \rangle$, i.e. the average dominance duration following a dominance of given duration T_i : the derivative of $\langle T_{i+1}|T_i \rangle$ evaluated at the average dominance provides an approximate value for the correlation coefficient [Schwalger, 2013].

We obtained the following result: in the alternation sequences obtained from our optimised model, pairs of successive dominances are structured differently when considering low-contrast stimuli or high-contrast stimuli. At low contrast, the sequence of dominance duration is close to that of a renewal process (**Fig. 5.16 A**), and dominance pairs are unstructured, as if they were generated by a renewal process. This is not the case at high contrast: we observe a deficit of short-long/long-short pairs of dominances, balancing an excess of short-short and long-long pairs of dominances (**Fig. 5.16 B**). This shows that dominances of similar durations tend to follow each other, which explains the correlations obtained for stronger stimuli. Remarkably, this prediction was confirmed in behavioural alternation series, where the same qualitative trend was observed (**Fig. 5.16 AB**, left). At high contrast, the slope of the conditional expectation $\langle T_{i+1}|T_i \rangle$ is clearly positive around the average dominance durations, in both sets of experimental and simulated data, indicating a significant and positive correlation.

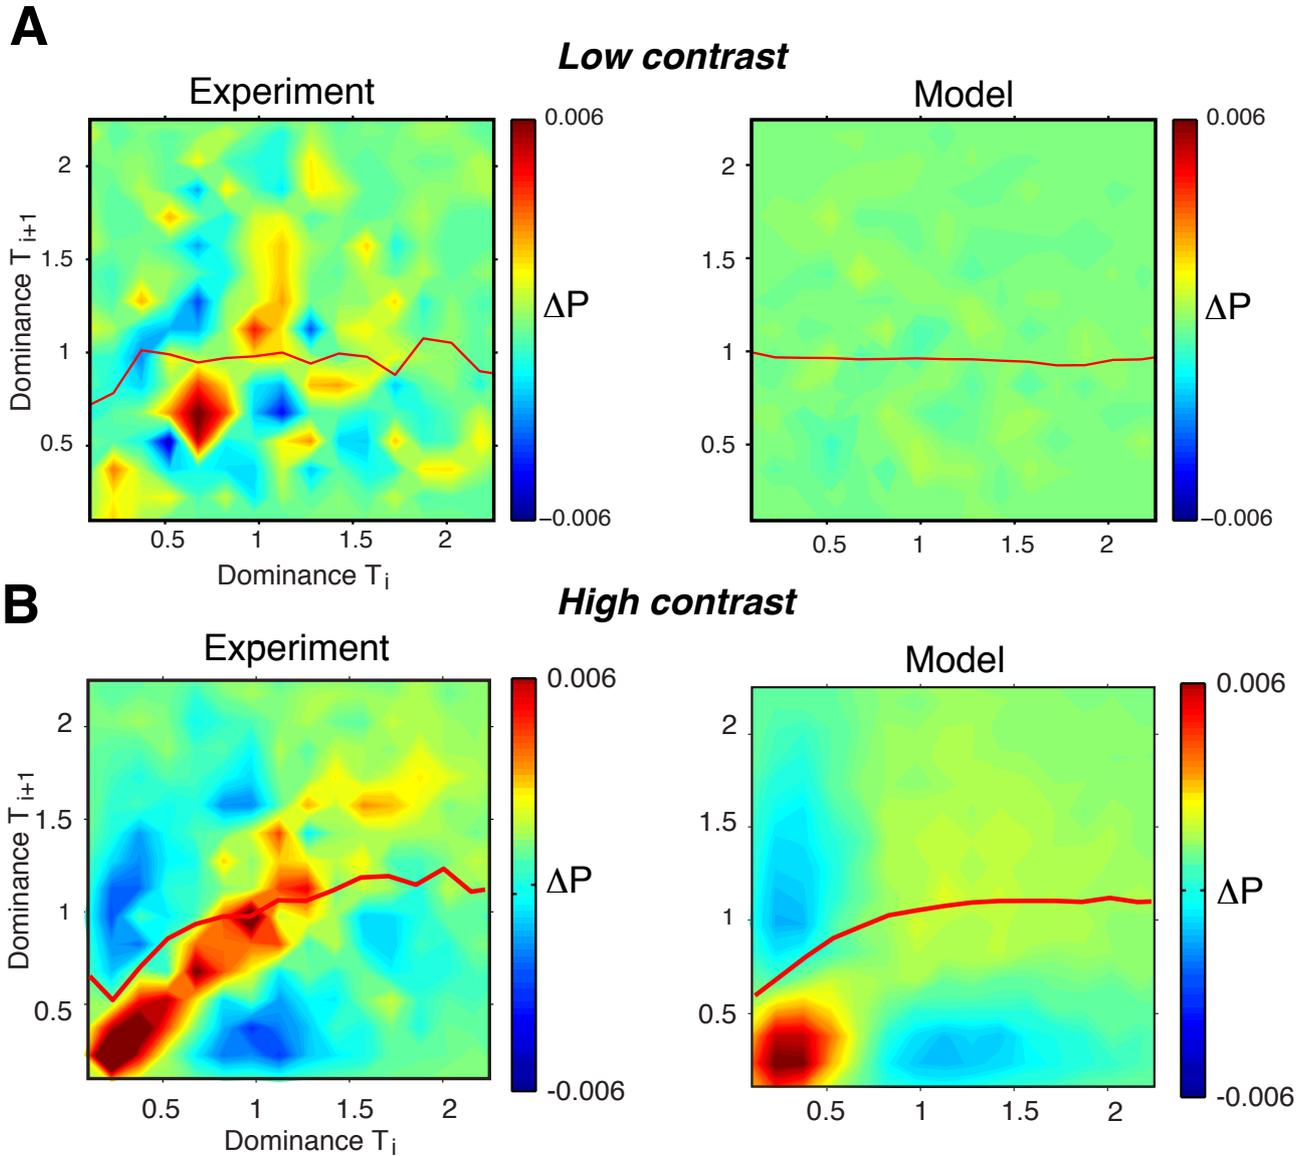


Figure 5.16: Correlations show a characteristic pattern for high stimulus contrast in both experimental (*left*) and simulated data (*right*). Excess joint probability ΔP (see (5.20)) of successive dominance durations T_i and T_{i+1} (contour plots) and conditional expectation $\langle T_{i+1}|T_i \rangle$ (red curves), in units of the mean dominance duration $\langle T \rangle$. **A** Correlation map for low stimulus contrast, $c_{dom} = c_{sup} < 0.125$ reflects that of a renewal process. **B** Correlation map for high stimulus contrast, $c_{dom} = c_{sup} > 0.5$ reflects that of a non-renewal process.

We next looked for further evidences of long-lasting memory effects at the different levels of the model's hierarchy, in an attempt to identify the dynamical mechanisms at the source of this peculiar correlation structure, and to determine the characteristic time-scales at which they operate (**Fig. 5.17**). For this we computed the mean expected dominance duration preceding and following each dominance period $\langle T_{i\pm j}|T_i \rangle$, this time for values of $j \geq 1$. In addition, to determine whether dominances of typically long or short duration contributed to non-renewal effects in the same manner, we sorted dominances durations in octiles, and computed mean expected durations for each octile separately (**Fig. 5.17 A**). We applied the same procedure at the evidence level, this time by observing the average activity $Z_i = \left(\frac{x+y}{2}\right)(t_i)$ in the evidence pool, *at reversal times* t_i , also sorted in the same octiles as previously (**Fig. 5.17 B**).

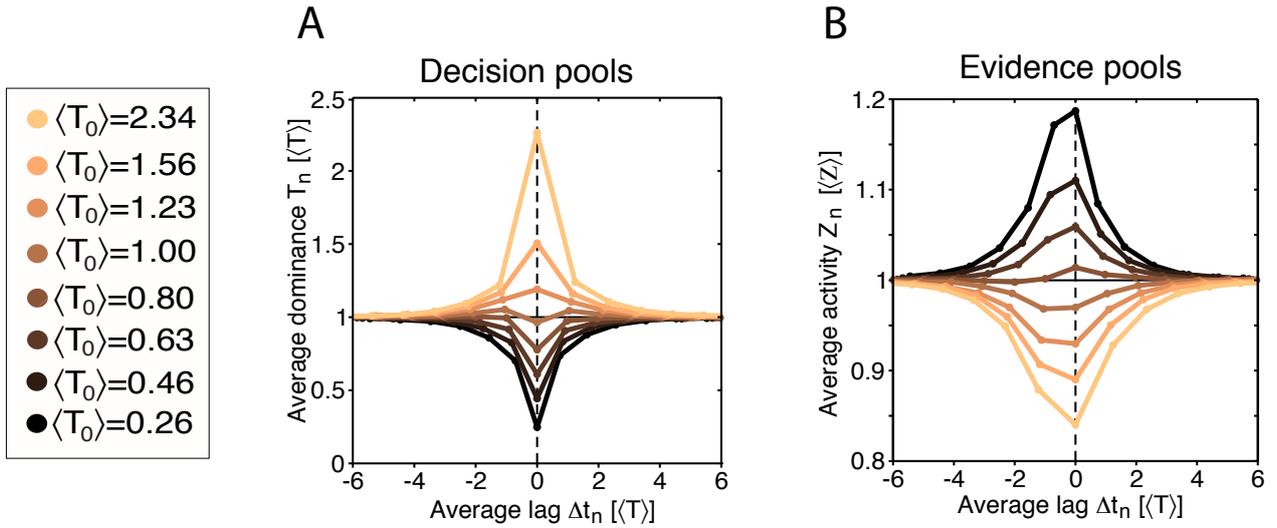


Figure 5.17: **A** Conditional expectation for successive dominance durations at the decision level, $\langle T_{i\pm j}|T_i \rangle$. **B** Conditional expectation for average evidence activity at the evidence level, $\langle Z_{i\pm j}|Z_i \rangle$ ($j \geq 1$). For high contrast levels, dominance periods are sorted into octiles, from shortest (beige) to longest (black), the averages of each cohort, and of the preceding and succeeding cohorts, are shown as a function of average lag time. Durations are in units of the overall average, $\langle T \rangle$.

These results show that, for extreme samples of the distribution of dominance durations, memory effects may range up to 3 times the average dominance duration, explaining the "clustered" structure of alternation series. Dominances of particu-

larly shorter (or longer) durations than the average could indeed be observed in consecutive sequences. Interestingly, a similar effect was observed in the evidence layer. Average activity was found to be the highest (resp. lowest) when measured at a reversal preceding the shortest (resp. longest) dominance durations. This effect was visible for up to 3 preceding and succeeding average dominances. This clearly pinpoints the origins of correlation in macroscopic alternation sequences to a slow, noisy drift of average activity at the level of sensory evidences. This is consistent with the idea that noise in slow adaptation currents can produce significant positive serial correlation in stochastic sequences, such as in multistable perception, or in the statistics of spiking neurons [Lehky, 1988, van Ee, 2009, Schwalger, 2013].

5.6.2 Slow fluctuating drift of sensory evidences

The correlation structure of alternation series arises naturally from the hierarchical structure of our model. We know from Section 5.5.4 that the dependence of the reversal threshold on average evidence activity induce modulations of average reversal times. Indeed, the average effective threshold $\langle \theta_{eff} \rangle$ decreases for increasing input strength, because the average activity in evidence populations $\langle Z \rangle$ increases, providing a qualitative explanation for L4. Importantly, within a given alternation series, where contrasts are fixed (here we assume contrasts are high, so that we obtain significant correlations), the effective reversal threshold also varies as $Z(t)$ varies.

Typically, higher values of Z_i (the value of $Z(t)$ at reversal times t_i) will yield smaller values of θ_{eff} : the decision level becomes increasingly sensitive when supported by larger activity at the evidence level (**Fig. 5.18 A**). This slow and fluctuating drift of average sensory evidences is precisely what produces this short-term history-dependence effect. While the evidence bias $\Delta(t)$ drives perceptual alternations, the average sensory activity $Z(t)$ fluctuates, slowly hovering around its average value $\langle Z \rangle$ (**Fig. 5.18 B**). The effective threshold varies locally, which translates into ‘bundles’ of dominances of comparable length. During trial segments where $Z(t) < \langle Z \rangle$, reversal thresholds will be consistently larger than average, and dominance periods, longer (**Fig. 5.18 C**). Conversely, during trial

segments where $Z(t) > \langle Z \rangle$, reversal thresholds will be consistently smaller than average, and dominance periods, shorter (**Fig. 5.18 D**).

However, since fluctuations of $Z(t)$ are also present for low and moderate contrast, and could therefore potentially yield similar threshold modulations, why do we not observe the same correlation between local fluctuations of average activity and local fluctuations of the threshold? The answer lies in how fast $Z(t)$ returns back to its average value over time. Recall that, in **Fig. 5.18 A**, we observe a significant correlation between activity and effective threshold, for values Z_i of $Z(t)$ measured *at reversal times* t_i . Between reversals, the average activity $Z(t)$ keeps fluctuating around its average value, typically relaxing towards it whenever deviations have occurred. Consider the case where $Z(t)$ has strongly deviated from its average, driven by the fluctuations of evidences. When dominance periods are short, as is the case at high contrast, and $Z(t) > \langle Z \rangle$, it may take several successive dominance for $Z(t)$ to reach its average again. In the meantime, for each reversal that occurs, Z_i is still be greater than average and the next dominance will also often be shorter than average. When dominance periods are long, however, as is the case for low and moderate contrasts, $Z(t)$ may be able to fully relax from deviations from its average value within a single dominance, and may even become smaller than $\langle Z \rangle$. Therefore a short dominance may not be systematically followed by a short dominance, but also by an average, or even a longer dominance. In other words correlations are suppressed when the modulation rate of average sensory activity is faster than the average reversal rate.

5.7 Summary

In this chapter, we have presented the main objective of this thesis, which was to propose a model of bistable perception able to reproduce the all aspects of the statistics of reversal times, and in particular, the scaling property. Adaptation-based models are unable to do so because the balance between adaptation and

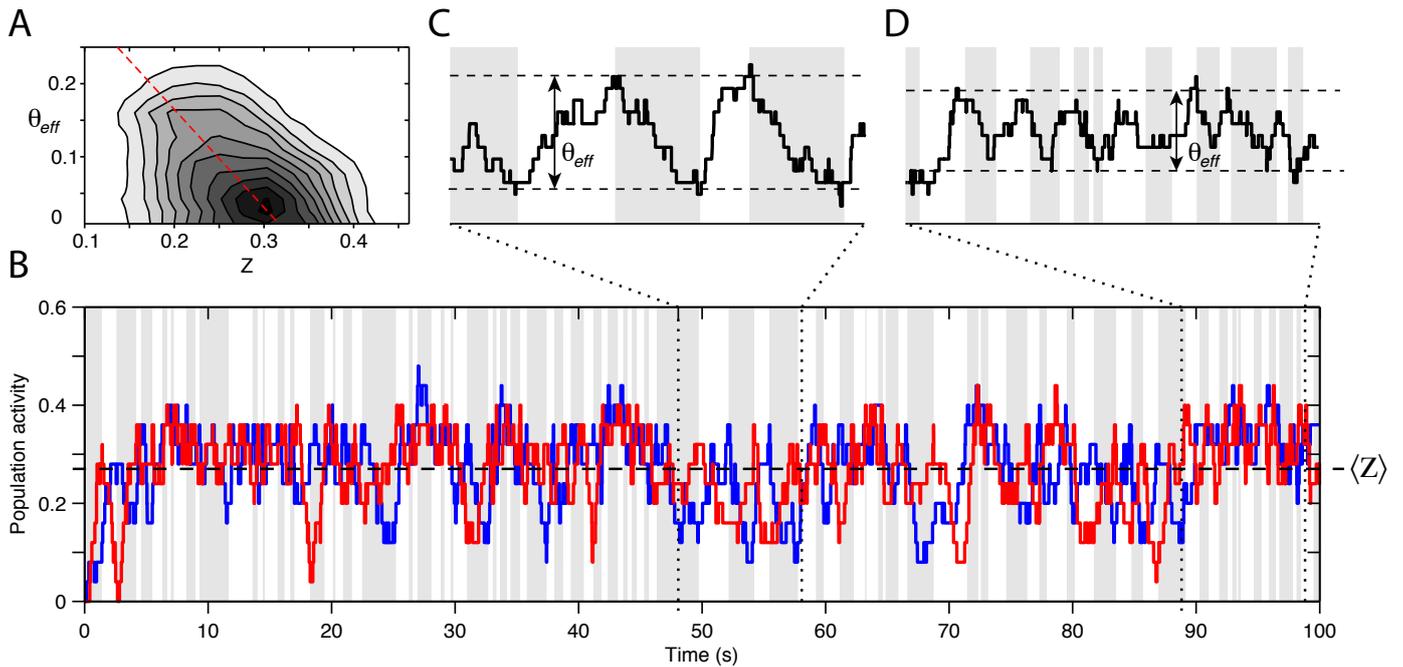


Figure 5.18: Our model explains the structure of observed correlations in terms of a slow drift of evidence activities. **A** Joint distribution of effective threshold θ_{eff} for differential evidence activity, revealing systematic dependence. **B**: representative simulation sequence with evidence activities, $x(t)$ and $y(t)$ (blue and red traces), of finite system ($N = 25$). Note slow drift of combined activity, $Z = x(t)/2 + y(t)/2$, with periods above- and below-average, $\langle Z \rangle$. **C** Differential activity (or evidence-bias) $\Delta(t) = x(t) - y(t)$ during a period of *lower* activity and *longer* dominance periods, due to *higher* θ_{eff} . **D** Differential activity during a period of *higher* activity and *shorter* dominance periods, due to *lower* θ_{eff}

noise they rely on is not robust to changes in stimulus strength. In our model, at the decision level, competition is implemented in the same manner via mutual inhibition. The difference is that we replaced self-adaptation by a stochastic integration at a lower level of representation (evidence level), which does not compromise the FPT statistics. Reversals are now driven by the accumulation of an evidence-bias $\Delta(t)$ to the perceptual threshold, i.e. the difference in activity between evidence populations supporting either competing percept. There are two independent components in the dynamics of the evidence-bias: the recovery of evidences supporting the suppressed percepts, and the habituation of evidences supporting the dominant percept. One important consequence of this *differential* integration process is that, unlike adaptation, its temporal properties can be modulated without affecting the

stability of perceptual representations at the decision level. This leaves the balance between deterministic and stochastic components of the dynamics unchanged, so that the scaling property is verified nearly independently of average dominance durations.

Our new model of multistable perception proposes detailed and comprehensive accounts for important statistical properties of the available behavioural data. In particular, it explains, both quantitatively and qualitatively, the peculiar input-dependence of average dominance durations, also known as Levelt's propositions. The scaling property, as well as the stereotypical distribution shape of distribution of dominance durations, are also obtained over most of the range of contrast combinations. Furthermore, the model reveals a specific structure of correlations in alternation series, and identify its origins to a simple mechanism: a slow and noisy modulation of sensory evidences. This prediction was confirmed in behavioural data, where the same qualitative contrast-dependence of correlations was observed. To conclude this study, we will discuss how the constrains imposed by behavioural data on the model's dynamics can be understood in the light of neurophysiological evidences and of other important aspects of perceptual decision-making.

6 Conclusions

Throughout this dissertation, I have presented the path of investigations which led to the proposal of a new model of multistable perception. The original objective of the study to reproduce scale-invariant reversal sequences, as observed in experiments, was successfully met. The resulting model provides a remarkable fit to the stationary statistics of dominance durations for the first three moments (mean, CV and skewness), for all combinations of stimulus strengths. At first-order, Levelt's propositions (under their most recent reformulation) are all well accounted for, both qualitatively and quantitatively. At second and third order, the scaling property is satisfied, and reversal time densities correspond to a Gamma distribution (with $c_v \approx 0.6$ and the skewness $\gamma_1 \approx 2c_v$), consistently with behavioural observations. To the best of my knowledge, this model is the first to reproduce the available experimental data to this level of detail. Previously introduced models only partially explain Levelt's propositions, usually for restricted combinations of stimulus strength. Importantly, they operate in a narrow noise-driven regime which prevents the scaling property from holding when average reversal rates vary, as is necessary to reproduce Levelt's propositions.

Additionally, our model carries several interesting implications beyond our original objectives. The model predicts non-stationary aspects of reversal sequences on a range of time-scales which had not been previously considered, partly because sequential correlations reported by experiments were generally considered

too weak. This lack of correlation has often justified the necessity for prior models to be confined to a noise-dominated regime. In contrast, our model operates in a strictly drift-dominated regime, where sequential correlations in alternation series arise from the non-stationary nature of the underlying dynamics. Perceptual alternations are thus not entirely memoryless, with effects of history-dependence sometimes lasting over several dominance periods. At high stimulus strength, alternation series become ‘bursty’, forming groups of successive dominance periods of comparable durations.

Our proposed framework relies on plausible neurophysiological assumptions which are consistent with observations reported in non-human primates. Furthermore, our model’s architecture and dynamics rediscover some important aspects of current models of decision-making. As such, it could constitute a suitable framework to study continuous detection and discrimination tasks in changing environments, where sensory inputs may reflect the simultaneous representation of multiple and subtly entangled stimuli.

Macroscopic laws reflect the nature of underlying processes

The unquestionable complexity of the brain’s structure and dynamics, along with the autonomous nature of its development, from single-cells to functional networks, is perhaps one of the most salient instance of self-organization in nature. Several aspects of brain dynamics are known to reproduce hallmarks of critical behaviour, such as multistability and scale-invariant fluctuations, which may reflect functional imperatives of neural systems [Beggs and Plenz, 2003, Petermann et al., 2009, Freyer et al., 2012, Hesse and Gross, 2014]. Similarly, multistable phenomena are thought to reveal the self-organized nature of perceptual representations, and reversal behaviour to manifest a self-induced instability ensuring that all plausible interpretations of the stimulus are explored [Friston et al., 2012, Pastukhov et al., 2013]. For systems operating at criticality, dynamical properties at different scales are singularly linked, in that the qualitative nature, rather than the details of their microscopic constituents, can directly reflect the emergence of macroscopic laws. Prior models have focused on detailing the mechanism underlying reversal

behaviour, and generally rely on mean-field assumptions which precludes them from interpreting macroscopic fluctuations as an emergent property of collective microscopic dynamics. In contrast, this thesis shows that the scaling property of multistable perception constrained the qualitative properties and dynamics of the underlying neural processes at different levels of the hierarchy of perceptual representation.

The first aspect of the scaling property, i.e. the invariance of shape of reversal times densities, is ensured by assuming spontaneous activity in metastable units at a *microscopic scale* (see Chapter 2-3). The second aspect of the scaling property, i.e. the specific Gamma-like shape of reversal times densities, requires stochastic accumulation at the *mesoscopic scale* to operate in a far-from-equilibrium and transient operating regime (see Chapter 4). Finally, at the *macroscopic scale*, or behavioural level, a specific structure of hierarchical interactions, featuring both feedforward and feedback projections was required to implement perceptual reversals, whilst preserving the higher-order statistics of alternation sequences (Chap. 5).

Microscopic scale, local attractors

At the very basis of our approach lies the idea that patterns of brain activity evoked by external stimulation correspond to the activation of ‘pre-existing’ states, and that such states can be also accessed spontaneously even in the absence of stimulation. We used stochastic bistable units to represent the dynamics of those states. Individually, these units can be thought of as local attractor networks, with which they share some fundamental properties: *local bistability*, *noise-driven transitions*, and *input-dependent* transition rates [Amit, 1995]. Attractor assemblies seem particularly well-suited to describe spontaneous brain activity: given adequately balanced recurrent excitation and inhibition, such networks become bistable may spontaneously transit between a ‘low’ or ‘high’ state of activity, driven by endogenous fluctuations. Transition rates between the two coexisting states typically depend on incoming synaptic activity and on the strength of recurrent connections [Amit

and Brunel, 1997].

The columnar organisation of cortical neurons, the structure of synaptic connections within cortical columns, and of lateral projections between neighbouring columns or clusters of columns, suggest a highly heterogeneous connectivity map dominated by local recurrent projections: a key feature of local attractor assemblies [Douglas et al., 1995, Douglas and Martin, 2007]. Although it is currently not known whether neural assemblies actually behave as stereotypically as our hypothetical metastable units do, evidences of sequential transitions between distinct modes of activity have been reported in recordings of cortical neurons in premotor, parietal and visual areas of primates [Abeles et al., 1995, Mattia et al., 2013, Latimer et al., 2015].

Mesoscopic and macroscopic scale, ‘nested attractors’

At the mesoscopic scale, when considering of a finite population of bistable units, we have seen that behavioural statistics impose strict conditions on the accumulation of activity leading to perceptual reversals. Although we have assumed that bistable units within the same pool do not interact, the energy landscape corresponding to the evolution of population activity is uneven, and the diffusion not necessarily homogeneous: in fact, this is only the case at and around steady-state activity. Since accumulation of activity must proceed far-from-equilibrium, population dynamics is constrained to a *transient* wandering in an *heterogeneous* and *granular* energy landscape. The implications of these characteristic features are perhaps more compelling when considering both evidence and decision level, to include the macroscopic scale at which perceptual reversals are observed.

In this context, slow transient dynamics emerges at the level of *evidence pools*, from successive stochastic hopping between local attractors. Accumulation of evidences ultimately drives transitions between global attractor states, observed in the form of perceptual reversals at the level of *decision pools*. This realises a ‘chaotic itinerancy’, in a ‘nested attractor’ landscape (i.e. ‘attractors within attractors’), during which a limited repertoire of accessible states is explored until a perceptual decision occurs, or in multistable perception, a perceptual reversal [Durstewitz and Deco, 2008, Gigante et al., 2009, Braun and Mattia, 2010, Friston

et al., 2012].

A similar type of slow, transient dynamics between distinct modes of activity has been implemented in networks of spiking-neurons, using an heterogeneous, clustered structure of connectivity forming a number of functional subnetworks [Litwin-Kumar and Doiron, 2012]. This is qualitatively equivalent to the result discussed in Section 3.4.3, we have proposed an extension of the Ehrenfest process which consider the collective dynamics of interacting bistable units, which can be thought of as an heterogeneous, multi-modular cortical network, or cluster of interacting cortical columns.

Patterns of interactions observed within cortical circuits are consistent with the picture proposed by clustered attractor networks. At the scales considered by such models, attractor-like states are observed in the form of spatially distributed and transient patterns of activity, rather than localised steady-states [Braun and Mattia, 2010, Wang, 2012]. Anatomical studies reveal similar ‘patchy’ patterns of connectivity in superficial layers of primates’ visual cortex [Lund et al., 2003, Tanigawa et al., 2005]. Remarkably, both spontaneous and evoked patterns of brain activity observed in imaging studies seem to directly reflect this peculiar organisation: while such activity is undeniably stochastic, it is far from random. The spatiotemporal structure of spontaneous neuronal responses is closely related to that of evoked responses; this led to the suggestion that spontaneous cortical dynamics rehearses reproducible patterns from a limited repertoire, autonomously revisiting states that may have been previously explored under sensory stimulation [Kenet et al., 2003, Beggs and Plenz, 2004, Han et al., 2008, Luczak et al., 2009, Berkes et al., 2011]. In a recent imaging study, transient brain activity during perceptual rivalry was reduced to sequential activation between a finite number of distributed stable activity states [Watanabe et al., 2014]. Reconstruction of the associated energy landscape revealed a ‘nested attractor’-like structure, in which exploration of intermediate local minima allowed transitions between more segregated, larger basins of attractions. Neurophysiological evidences therefore appear to support a ‘nested attractor’ framework to describe and for which this thesis proposed an effective implementation.

Top-down influences balance ‘Exploration’ and ‘Exploitation’

A growing body of evidence supports the idea that multistability arises from interplay between neural processes at different levels of the hierarchy of visual processing, and involve mutual interactions between both low- and high-level representations of the stimulus’ properties. [Leopold and Logothetis, 1999, Blake and Logothetis, 2002, Sterzer et al., 2009, Kornmeier et al., 2009, Kornmeier and Bach, 2012]. The ineluctability and inherent stochasticity of reversal behaviour may reflect a perceptual analogue of the ‘exploration-exploitation dilemma’, in that stochastic exploration adjusts the relative contribution of prior knowledge and current sensory input [Leopold and Logothetis, 1999, Hohwy et al., 2008, Sundareswara and Schrater, 2008, Moreno-Bote and Knill, 2011]. Perception may operate in a marginally stable regime to reach an optimal compromise between perceptual stability and sensitivity, to ensure a continued representation of the stimulus while allowing alternative interpretations to be considered [Friston et al., 2012, Pastukhov et al., 2013].

For prior models relying on a self-adaptation mechanism to implement perceptual reversals, this specific operating point typically lies in a narrow parameter range in the vicinity of an oscillatory bifurcation, with the important drawback that the scaling property can not be satisfied (see Section 5.2). Comparatively, our hierarchical model expresses this balance between exploration and exploitation in a more robust manner. On the one hand, conscious perceptual dynamics is very sensitive to changes in the subconscious dynamics of sensory evidences, perceptual thresholds are small, and only a few activations (or inactivations) are needed to initiate perceptual reversals. On the other hand, dominant appearances are not easily destabilised by endogenous noise, and remain stable between successive reversals.

In our proposed framework, the regime in which reversal behaviour occurs reveals a joint and coordinated hierarchical dynamics between evidence-driven influences and decision-driven influences. While initial percept selection directly results from feedforward excitation, reversals result from a combination of feedforward

excitation, and feedback inhibition. Between successive reversals, sensory dynamics consists of two separate and antiphasic components: evidences supporting the currently suppressed percept recover, while evidences supporting the alternative (dominant) percept habituate, driven by the current perceptual state. In that sense, the resulting dynamics highlights an ‘exploration’/‘exploitation’ balance, with top-down influences playing a central, albeit not exclusive role in enabling exploratory behaviour. Changes in perception at a higher-level of cognition prompt the re-organisation of earlier sensory areas: while the dominant appearance is steadily maintained, a stochastic exploration is carried out at the lower-level [Leopold and Logothetis, 1999]. The idea that sensory information supporting alternative or dominant appearances is continuously integrated and discarded (‘explained away’), has previously been highlighted by several authors in a comparable way [Sundareswara and Schrater, 2008, Hohwy et al., 2008].

Whether perceptual reversals result from bottom-up ‘passive’ neural processes such as self-adaptation, or from top-down ‘active’ cognitive processes such as attention, is not definitely known, and remains an active topic of research and debate. Evidences in support of both alternative have been reported, and some experimental studies suggest that they may not mutually exclusive and could cooperatively contribute to reversal dynamics [Lumer and Rees, 1999, Kornmeier et al., 2009, Wang et al., 2013]. Interestingly, several authors have formulated qualitative mechanisms to incorporate both bottom-up and top-down influences in an attractor-like picture, in which the energy landscape associated with perceptual choices is continuously reshaped by both unconscious sensory integration and conscious experience [Friston, 2003, Hohwy et al., 2008, Friston, 2010, Friston et al., 2012, Kornmeier and Bach, 2012, Wang, 2012]. Our proposed framework explicitly materializes this scenario: the profile of the energy landscape associated with accumulation of sensory evidence between reversals is not static, rather, it adapts with incoming sensory input, via bottom-up influences, and conscious experience, via top-down influences.

It is interesting to note that our model offers an alternative perspective on the exploration-exploitation dilemma, as originally formulated in reinforcement learning, and as practically implemented in Bayesian models of multistable perception:

it does not assume exploratory behaviour to reflect a stochastic rule, such as neural sampling [Sutton and Barto, 1998, Sundareswara and Schrater, 2008, Moreno-Bote and Knill, 2011]. Here, exploration emerges as a natural by-product of our model’s hierarchical and reciprocal structure of interactions. This suggests that, in contrast to the passive self-adaptation mechanism used in previous models, our active, percept-driven habituation mechanism may be interpreted in terms of an inferential process, such as predictive coding [Hohwy et al., 2008].

Relation to probabilistic inference by sampling

Fluctuations in neural activity may serve an important functional role in the neural computations underlying perceptual decisions. In agreement with Helmholtz’s seminal intuition that visual perception is an inference process, there is mounting evidence that such variability can enable neural circuits to perform perceptual decisions probabilistically, by representing sensory uncertainty in the form of probability distributions [von Helmholtz, 1866, Hoyer and Hyvärinen, 2003, Knill and Pouget, 2004]. The way such mechanisms are effectively achieved within neural assemblies constitute an increasingly active area of research [Hoyer and Hyvärinen, 2003, Fiser et al., 2010, Haefner et al., 2016, Orbán et al., 2016]. Fluctuations in the firing activity of a single-neuron may directly reflect an inherent uncertainty of the quantity or object it represents [Deneve, 2008a,b, 2012]. In probabilistic population codes, fluctuations in individual firing activity over a whole neural population can represent uncertainty in terms of a complete probability distribution [Ma et al., 2006, Beck et al., 2008, Pouget et al., 2013]. There is also mounting evidence that neural circuits implements approximate probabilistic inference by means of stochastic sampling [Hoyer and Hyvärinen, 2003, Fiser et al., 2010, Moreno-Bote and Knill, 2011, Gershman et al., 2012, Haefner et al., 2016, Orbán et al., 2016].

The present study highlights a situation in which the transient nature of nonequilibrium fluctuations are of particular importance to account for important aspects of behavioural data (see Chapters 4). This is qualitatively different from perceptual transitions in adaptation-based models, in which reversals are generally

noise-dominated, and reflect fluctuations around equilibrium. These dynamical features are not exclusive to multistable perception, and have already been discussed in the broader context of perceptual decision-making. It has been proposed that stochastic transient dynamics across metastable attractors may reflect neural processes involved in cognition more realistically than classical fixed-point attractor states [Durstewitz and Deco, 2008, Braun and Mattia, 2010, Friston et al., 2012, Wang, 2012]. Given the inherent non-stationary nature of cortical activity, that the reproduction of important statistical features of multistable perception was only possible in this regime, may not be incidental: far-from-equilibrium operation may indeed reflect a functional feature of neural representations underlying perceptual dynamics.

It has long been understood that Bayesian inference can be approximated by stochastic systems assembled from discrete, binary components (variantes of Boltzmann machines; [Hinton and Sejnowski, 1986], Learning and relearning in Boltzmann machines; Markov random fields; [Robert and Casella, 1998, Hinton, 2014]), not unlike the discrete, bistable units postulated here. It has also been pointed out that such systems would naturally exhibit multi-stability, as well as other top-down effects during perceptual inference [Gershman et al., 2009, 2012, Hinton, 2014]. However, these models are of limited practical value, because they converge very slowly, The reason is that the dynamics operates near equilibrium (Glauber dynamics), where the state space is extremely large.

In this context, E. T. Jayne's information-theoretic approach to statistical physics appears to have important consequences, for his approach can be applied to study fluctuations in discrete stochastic systems (which are capable of approximating statistical inference) by extending the principle of entropy maximization to the path entropy of microtrajectories [Jaynes, 1983, Ghosh et al., 2006, Pressé et al., 2013]. A key concept is potency (or caliber), which measures the fraction of all microtrajectories that lead to a substantial change in the systems macrostate. The two aspects of greatest interest in the present context are that (i) potency grows with distance from equilibrium and (ii) potency decreases with system size. Thus, the

principles of maximum caliber imply that fluctuations in a smaller system, operating farther from equilibrium, will shift macrostates faster and more reliably than fluctuations in larger system near equilibrium. In other words, there may be compelling physical reasons for approximating perceptual inference with a dynamical system that operates (i) far from equilibrium and (ii) with a small population of discrete units, both of which are important requirements for our model to reproduce the statistics of multistable perception.

Inter-individual disparities

An additional interesting implication of our model concerns the inter-individual disparity of reversal rates, a long-known and remarkable feature of perceptual rivalry. For the same rivalry display, average reversal rates between subjects can span over an order of magnitude [Fox and Herrmann, 1967, Walker, 1975]. The reasons for such differences are not well understood. Recent imaging studies have proposed that such inter-individual differences in perceptual timing may in fact reflect anatomical differences: grey matter density in brain areas mediating perceptual rivalry (notably in the superior parietal lobule) was shown to significantly correlate with reversal rate [Kanai et al., 2011, Kleinschmidt et al., 2012, Watanabe et al., 2014]. These results have usually been interpreted in terms of the strength and nature of connectivity between these areas and earlier sensory areas [Kleinschmidt et al., 2012, Megumi et al., 2015].

Another study reports strong correlations between concentrations of γ -Aminobutyric acid (GABA) - the principal inhibitory neurotransmitter - in visual areas, and average reversal times for three types of multistable displays (BR, MIB and KDE) Van Loon et al. [2013]. The authors were able to reproduce these results in an adaptation-based model by modulating the strength of mutual-inhibition.

Here we contribute to this debate by proposing an alternative - and compatible - perspective on this issue. Note that, in the GE process, the characteristic time of evidence pools can be set to arbitrarily small or large values, without compro-

mising its overall dynamics or the statistics of alternation sequences. Differences in the individual activation rates of bistable units, which intrinsically define the characteristic time-scale of collective activation, could account for the wide differences in average reversal rate reported between observers.

There are several ways in which this time-constant could be modulated. First, since bistable units can be thought of as a recurrently coupled network of neurons, different degrees of self-excitation would yield shallower or deeper basins of attraction, and thereby faster or slower individual transition rates, respectively. Second, the level of interactions between bistable nodes has a direct effect on the characteristic time-scale at which the evaluation of sensory evidence operates, whilst preserving higher-order properties (as discussed in Section 3.4.3). Finally, recall that, during the habituation phase, when evidence supporting the dominant percept is gradually suppressed, the characteristic time-scale of habituation would also be affected by the weight of top-down inhibition: a stronger coupling could in principle reduce the average dominance period (as also suggested in [Kleinschmidt et al., 2012]).

Limitations and outlook

To capture more qualitative aspects of multistable perception, our approach remained intentionally simplistic, and therefore may not reproduce all details of more elaborate experiments. For instance, events occurring on a much shorter characteristic time-scale than that of the individual bistable units can not be well captured in the collective dynamics. This prevents the model, under its present form, to account for experimental results involving fast transient stimulus. One solution could be to include different types of bistable units at the evidence level, some with fast switching rates, others with low switching rates, or to consider inhomogeneous interactions, similarly to clustered attractor networks [Litwin-Kumar and Doiron, 2012].

Conversely, particularly slow modulations such as critical slowing can not be readily obtained [Leopold and Maier, 2012]. Here, the solution would require slow memory pools to be introduced, as previously proposed in a previous model of multistability [Gigante et al., 2009].

There is also no explicit spatial dimension in the model. Experiments involving spatially extended stimuli, for which perceptual grouping may play a role and produce effects such as ‘piecemeal rivalry’ and ‘travelling waves’ are thus out of the model’s scope. These issues have already been investigated in traditional energy-well models, by considering that ‘complete’ rivalry of large stimuli could be represented by the interaction between ‘partial’ or ‘local’ rivalries responding to different regions of the stimulus [Kang, 2009]. A similar extension could also be considered for the present model.

Our approach sought to generalise the picture offered by drift-diffusion models, and provide a neurophysiological basis for the stochastic decision variables thought to underlie perceptual decisions. By reconciling the dynamics of neural attractor assemblies and of stochastic diffusion-to-bound within a hierarchical structure, our practical implementation of a ‘nested attractor’ confirmed the promising nature of such framework, as predicted by several authors [Gigante et al., 2009, Braun and Mattia, 2010]. Gratifyingly, our model expresses important features of decision-making theories and may carry important implications beyond the context of multistable perception.

Presumably, visual interpretation of natural scenes involves a general disambiguation mechanism, which becomes increasingly engaged as the uncertainty about the true nature of the stimulus increases [Yuille and Kersten, 2006, Sterzer and Rees, 2008, Wang et al., 2013, Pastukhov et al., 2013]. Multistable perception can be thought of as being positioned at one end of this ‘spectrum’ of increasing stimulus uncertainty, in that it reflects a perfectly ambiguous scenario where the action of this mechanism phenomenally affects perceptual dynamics, making perceptual reversals particularly noticeable. However, multistability may be a widespread feature of perceptual dynamics and operate unnoticed in many situations, for instance, when considering real-world stimuli where competition between multiple subtly ambiguous elements would not trigger reportable conscious changes [Deco et al., 2007, Pastukhov et al., 2013]. Therefore, neural processes un-

derlying multistable perception may operate on the basis than the disambiguation mechanism which allows correct categorization of all visual inputs.

The success of the present study in accounting for empirical observations of multistable phenomena, in a mechanically simple and neurophysiologically plausible manner, suggests that we may have, at least in parts, uncovered the dynamical mechanisms and neural circuits of perceptual inference. We therefore expect that our models architecture and dynamics could be extended to perform perceptual decisions in increasingly complex and ecologically realistic scenarios.

At this point, I see two main research directions to further develop of this framework. A first direction would be to consider continuous classification tasks in changing environments and a second direction would consider complex, high-dimensional stimuli, with a large number of relevant features. Both directions could eventually be combined into a comprehensive theory of normal vision.

Consider an environment with two alternating stimuli appearing and disappearing at random times. In this simple case, the identity of the current stimulus can be inferred optimally, for example by means of a sequential probability ratio test (SPRT) [Wald, 1947]. In more challenging situations with multiple stimulus alternatives and/or stimulus statistics that change over time, the computational demands of optimal inference grow prohibitively [Moran, 2015]. Even under such conditions, optimal inference can be approximated by stochastic dynamical systems that incorporate either non-linear interactions to satisfy additional time-dependent constraints [Deneve, 2012, Veliz-Cuba et al., 2016] or time-dependent perceptual threshold (see for instance [Hanks et al., 2011, Drugowitsch et al., 2012]). In view of the close resemblance between these systems and the model proposed here, it is evident that the present model could readily be adapted to performing continuous inference in a changing environment. For example, one adaptation beneficial for inference could be to match the switching rates of evidence units to the presentation times of the stimuli in question, and another one could be to employ feedback inhibition to suppress (explain away) evidence for already discriminated stimuli, such as to increase sensitivity for other stimuli.

Generalizing optimal inference to high-dimensional stimuli with multiple relevant features is mathematically challenging and the equivalence of stochastic dynamical models has not been established [Bitzer et al., 2014, Veliz-Cuba et al., 2016]. However, multiple-choice discriminations can be realized by coupling several diffusion models (each accumulating evidence for a different choice) by means of hierarchical interactions [Mazurek et al., 2003, Bogacz et al., 2006, Ditterich, 2010, Bollimunta and Ditterich, 2012]. Similar extensions could be envisaged for our model, creating a hierarchical architecture very much akin to a global neuronal workspace [Dehaene and Changeux, 2011]: unconscious sensory nodes projecting onto higher-level assemblies which, in turn, relay decisional outcomes back to sensory nodes.

The model we have introduced in this thesis, although introduced to solve the particular case of multistable perception, appears adaptable and generalizable in ways which are compatible with recent advances in theories of perceptual decision making, and . The question of whether our proposed framework can be successfully transposed to more general situations, and ultimately, to normal vision, offers truly exciting future research prospects.

Appendices

General solution of the Ehrenfest process

We derive the formula for $P_{ik}(t)$, the probability to start with i active nodes at time $t = 0$ and have k active nodes at time t . For this, we start from the probability generating function ((3.18)):

$$G_i(z, t) = (\lambda z + \mu)^N \left[1 + \lambda \frac{1-z}{\lambda z + \mu} e^{-t/\tau} \right]^{N-i} \left[1 - \mu \frac{1-z}{\lambda z + \mu} e^{-t/\tau} \right]^i \quad (\text{A.1})$$

To expand this in powers of z and identify $P_{ik}(t)$, we use the generating function of Krawtchouk polynomials:

$$\left(1 - \frac{\mu}{\lambda} x \right)^i (1+x)^{N-i} = \sum_{n=0}^N \binom{N}{n} K_n(i) x^n \quad (\text{A.2})$$

We recall:

$$K_n(x) = {}_2F_1(-n, -x, -N, 1/\lambda) = \sum_{k=0}^n \frac{(-n)_k (-x)_k}{(-N)_k k!} \left(\frac{1}{\lambda} \right)^k \quad (\text{A.3})$$

$$\begin{aligned}
G_i(z, t) &= \sum_{n=0}^N \binom{N}{n} K_n(i) \lambda^n \mu^{N-n} (1-z)^n \left(1 + \frac{\lambda}{\mu} z\right)^{N-n} e^{-\frac{nt}{\tau}} \\
&= \sum_{n=0}^N \binom{N}{n} K_n(i) \lambda^n \mu^{N-n} \sum_{k=0}^N \binom{N}{k} \frac{\lambda^k}{\mu} K_k(n) e^{-\frac{nt}{\tau}} z^k \\
&= \sum_{k=0}^N \left[\binom{N}{k} \frac{\lambda^k}{\mu} \sum_{n=0}^N \binom{N}{n} \lambda^n \mu^{N-n} K_i(n) K_k(n) e^{-\frac{nt}{\tau}} \right] z^k
\end{aligned}$$

Where we have used expression (A.2) twice, and also the identity $K_n(i) = K_i(n)$.

We can identify the probability $P_{ik}(t)$ as:

$$P_{ik}(t) = \binom{N}{k} \frac{\lambda^k}{\mu} \sum_{n=0}^N \binom{N}{n} \lambda^n \mu^{N-n} K_i(n) K_k(n) e^{-\frac{nt}{\tau}} \quad (\text{A.4})$$

Details on FPT moments derivation

To calculate the mean first-passage-time, we start from expression (3.32). At first-order, we have

$$\frac{d}{dx} \left(\frac{K_i(x)}{K_j(x)} \right) \Big|_{x=0} = \frac{K_i'(x)K_j(x) - K_j'(x)K_i(x)}{K_j^2(x)} \Big|_{x=0} = [K_i'(x) - K_j'(x)] \Big|_{x=0} \quad (\text{A.5})$$

We now need to calculate the derivatives of the Krawtchouk polynomials with respect to x , which actually reduces to calculating the derivative of the pochhammer symbol. For this, we introduce the first-order derivative of the Gamma function, also known as the Digamma function:

$$\psi_0(x+1) = \psi_0(x) + \frac{1}{x} \quad (\text{A.6})$$

By using the recurrence relation defined by (A.6), we can evaluate the derivative

of the Pochhammer symbol (3.26) in zero as follows:

$$\begin{aligned}
\frac{d}{dx}(x)_k &= (x)_k(\psi_0(x+k) - \psi_0(x)) = (x)_k \cdot \sum_{r=0}^{k-1} \frac{1}{x+r} \\
&= (x)_k \left(\frac{(x+1)\dots(x+k-1)}{x(x+1)\dots(x+k-1)} + \dots + \frac{x(x+1)\dots(x+k-2)}{x(x+1)\dots(x+k-1)} \right) \\
&= \underbrace{(x+1)\dots(x+k-1)}_{\xrightarrow{x \rightarrow 0} (k-1)!} + \dots + \underbrace{x(x+1)\dots(x+k-2)}_{\xrightarrow{x \rightarrow 0} 0}
\end{aligned} \tag{A.7}$$

This takes the simple form:

$$\frac{d}{dx}(x)_k \Big|_{x=0} = \begin{cases} (k-1)! & \text{if } k > 0, \\ 0 & \text{if } k = 0. \end{cases} \tag{A.8}$$

Similarly, at second order:

$$\frac{d^2}{dx^2} \left(\frac{K_i(x)}{K_j(x)} \right) \Big|_{x=0} = [K_i''(x) - K_j''(x) - 2(K_i'(x) - K_j'(x))K_j'(x)] \Big|_{x=0} \tag{A.9}$$

For the Pochhammer symbol:

$$\begin{aligned}
\frac{d^2}{dx^2}(x)_k &= (x)_k[\psi_0(x+k) - \psi_0(x)]^2 + (x)_k[\psi_1(x+k) - \psi_1(x)] \\
&= (x)_k \left\{ \left(\sum_{r=0}^{k-1} \frac{1}{x+r} \right)^2 - \sum_{r=0}^{k-1} \left(\frac{1}{x+r} \right)^2 \right\} \\
&= 2(x)_k \sum_{\substack{r,r'=0 \\ r < r'}}^{k-1} \frac{1}{x+r} \frac{1}{x+r'}
\end{aligned}$$

As before, we can explicit further explicit the sum:

$$\begin{aligned} \sum_{r < r'} \frac{1}{x+r} \frac{1}{x+r'} &= \left[\frac{1}{x} \frac{1}{x+1} + \frac{1}{x} \frac{1}{x+2} + \dots \right. \\ &\quad + \frac{1}{x+1} \frac{1}{x+2} + \frac{1}{x+1} \frac{1}{x+3} + \dots \\ &\quad \left. + \frac{1}{x+2} \frac{1}{x+3} + \frac{1}{x+2} \frac{1}{x+4} + \dots \right] \\ &= \left[\underbrace{(x+2)(x+3) \dots (x+k-1) + (x+1)(x+3) \dots (x+k-1) + \dots}_{\text{under-braced expression}} \right. \\ &\quad + x(x+3) \dots (x+k-1) + x(x+2)(x+4) \dots (x+k-1) + \dots \\ &\quad \left. + \dots \right] \frac{1}{(x)_k} \end{aligned}$$

Again, most terms will vanish when taking the limit $x \rightarrow 0$, except the under-braced expression We introduce the harmonic number $H_n = \sum_{k=1}^n k^{-1}$, and finally obtain

$$\begin{aligned} \frac{d^2}{dx^2}(x)_k &= 2(k-1)! \left[1 + \frac{1}{2} + \dots + \frac{1}{k-1} \right] \\ &= \begin{cases} 2(k-1)!H_{k-1} & \text{if } k > 0, \\ 0 & \text{if } k = 0. \end{cases} \end{aligned} \tag{A.10}$$

The same method was used to obtain the two following orders.

General solution of the balanced Poisson process

We start from the generating distribution (4.8), which we expand in powers of z using the binomial formula:

$$\begin{aligned}
G(z, t) &= e^{-\frac{t}{\tau}} e^{\frac{t}{\tau}(\lambda z + \frac{\mu}{z})} \\
&= e^{-\frac{t}{\tau}} \left[\sum_k \frac{1}{k!} \left(\frac{t\lambda z}{\tau} \right)^k z^k \right] \left[\sum_{k'} \frac{1}{k'!} \left(\frac{t\mu}{\tau} \right)^{k'} z^{-k'} \right] \\
&= e^{-\frac{t}{\tau}} \sum_n \sum_k \left(\frac{t}{\tau} \right)^{2k+n} \frac{(\sqrt{\lambda\mu})^{2k}}{k!(k+n)!} (\lambda\mu)^{\frac{n}{2}} \left(\frac{\lambda}{\mu} \right)^{\frac{n}{2}} z^n \\
&= \sum_n \left\{ \left(\frac{\lambda}{\mu} \right)^{\frac{n}{2}} e^{-\frac{t}{\tau}} \sum_k \frac{1}{k!(k+n)!} \left(2\frac{t}{\tau} \sqrt{\lambda\mu} \right)^{2k+n} \right\} z^n \\
&= \sum_n \left\{ \left(\frac{\lambda}{\mu} \right)^{\frac{n}{2}} e^{-\frac{t}{\tau}} I_n \left(2\frac{t}{\tau} \sqrt{\lambda\mu} \right) \right\} z^n \tag{A.11}
\end{aligned}$$

Due to the negative powers in z , each coefficient in the power expansion is defined by an infinite series which can be identified to the modified Bessel function of the first kind $I_n(x)$.

Equilibrium distribution of the CIR process

After a change of variable $y = \lambda(1-x) + \mu x$ in the Fokker-Planck equation for the CIR process (4.18):

$$\tau \frac{\partial p(y, t)}{\partial t} = -\frac{dy}{dx} \frac{\partial}{\partial y} \left[\left(\lambda - \frac{\lambda - y}{\lambda - \mu} \right) p(y) \right] + \frac{1}{2N} \left(\frac{dy}{dx} \right)^2 \frac{\partial^2}{\partial y^2} [yp(y)] \tag{A.12}$$

$$= -\frac{\partial}{\partial y} [(2\lambda\mu - y)p(y)] + \frac{1}{2} \frac{\partial^2}{\partial y^2} \left[\frac{(\lambda - \mu)^2}{N} yp(y) \right] \tag{A.13}$$

We obtain the stationary solution p_{st} by integrating (A.13) with the time-derivative set to zero:

$$p_{st} = C \cdot y^{\frac{4\lambda}{\sigma^2} - 1} e^{-\frac{2y}{\sigma^2}} \tag{A.14}$$

Where $\sigma^2 = \frac{(\lambda - \mu)^2}{N}$, and $C = \left[\Gamma\left(\frac{4\lambda}{\sigma^2}\right) \frac{\sigma^2}{2} \frac{4\lambda}{\sigma^2} \right]^{-1}$ an integration constant we determine by normalisation of p_{st} . After reverting the change of variable, we obtain:

$$p_{st} = \frac{1}{\Gamma\left(\frac{4\lambda}{\sigma^2}\right) \frac{\sigma^2}{2} \frac{4\lambda}{\sigma^2}} (\lambda - (\lambda - \mu)x)^{\frac{4\lambda}{\sigma^2} - 1} e^{-\frac{2}{\sigma^2}(\lambda - (\lambda - \mu)x)} \tag{A.15}$$

We can compare this distribution to the binomial distribution of parameter λ , the stationary distribution for the Ehrenfest process (3.6):

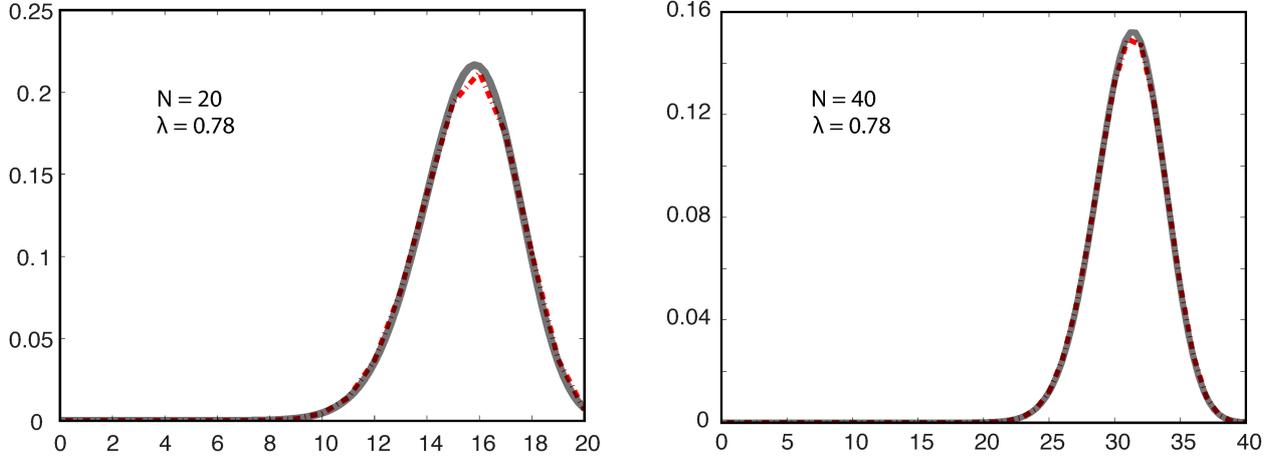


Figure A.1: Comparison of the stationary distributions for the CIR (solid grey line) and the Ehrenfest process (dotted red line), with $\lambda = 0.78$ and $\mu = 1 - \lambda$, with $N = 20$ (left) and $N = 40$ (right).

Alternative FPT moments formulae (numerical)

Here, we detail the method to obtain the moments of the first-passage-time density for the Fokker-planck equation with drift $\mu(x, x_{in})$ and diffusion $\sigma^2(x, x_{in})$:

$$\frac{\partial p(x, t)}{\partial t} = -\frac{\partial}{\partial x} (\mu(x, x_{in}) p(x, t)) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (\sigma^2(x, x_{in}) p(x, t)) \quad (\text{A.16})$$

Starting with an initial condition x_0 , we want to know the time $T(x_0)$ for the system to reach a certain threshold θ . $T(x_0)$ is a random variable with probability distribution $g(t, x_0)$.

The density $g(t, x_0)$ satisfies the associated backward Fokker-Planck equation:

$$\frac{\partial g(t, x_0)}{\partial t} = +\mu(x_0, x_{in}) \frac{\partial}{\partial x_0} g(t, x_0) + \sigma^2(x_0, x_{in}) \frac{\partial^2}{\partial x_0^2} g(t, x_0) \quad (\text{A.17})$$

Defining the n -th moment of random variable $T(x_0)$ as: $T_n(x_0) = \int_0^\infty t^n g(t, x_0) dt$, we can obtain the mean and variance of the first-passage time by solving the following set of equation (using the fact that $T_0 = 1$ by normalisation of distribution g) [Tuckwell, 1988/2008]. We can therefore obtain moments

recursively:

$$\begin{aligned} -\frac{1}{\sigma^2(x_0, x_{in})} &= \frac{\mu(x_0, x_{in})}{\sigma^2(x_0, x_{in})} T_1'(x_0) + T_1''(x_0) \\ -2\frac{T_1(x_0)}{\sigma^2(x_0, x_{in})} &= \frac{\mu(x_0, x_{in})}{\sigma^2(x_0, x_{in})} T_2'(x_0) + T_2''(x_0) \end{aligned}$$

Setting $\psi(x) = \int_{\infty}^x dy \frac{\mu(x_0, x_{in})}{\sigma^2(x_0, x_{in})}$ leads to the following integral formulae, which can be evaluated numerically:

$$\begin{aligned} T_1(x_0) &= \int_{x_0}^{\theta} dx \left[\int_0^x dy \frac{e^{\psi(y)}}{\sigma^2(x_0, x_{in})} \right] e^{-\psi(x)} \\ T_2(x_0) &= n \int_{x_0}^{\theta} dx \left[\int_0^x dy \frac{T_1(y)}{\sigma^2(x_0, x_{in})} e^{\psi(y)} \right] e^{-\psi(x)} \end{aligned}$$

Optimized parameters

Optimized model parameters are summarized here.

τ_{int}	τ_{dec}	$W_{E,dec}$	$W_{I,dec}$	$W_{E,ff}$	$W_{I,ff}$
1.9494	0.017669	15.2053	33.3775	152.1868	32.1033
$W_{I,fb}$	θ_{dec}	α	β	γ	
2.3402	-4.9383	0.65552	0.082017	0.070875	

Vicinity of optimal parameters set

To confirm that the fit with experimental data obtained from the minimization procedure was indeed optimal, and could not be further improved, we studied the behaviour of the fit error in the vicinity of the optimal parameter set.

For each parameter α_i , 10 values $\alpha_i^{(j)}$ in are picked in the direct vicinity of the optimal parameter α_i^{opt} . For each $\alpha_i^{(j)}$, the fit error $\bar{R}^{(j)}$ is computed, defining a scatter plot of the pairs $(\alpha_i^{(j)}, \bar{R}^{(j)})$ which is approximated by a quadratic function. For most parameters the estimated quadratic function is convex, so that the corresponding coefficient of the Hessian matrix associated with the fit error is positive. Also, the estimated extremum for each parabola is close to the corresponding optimal parameter. These results indicate that this parameter set is indeed a minimum for the fit error.

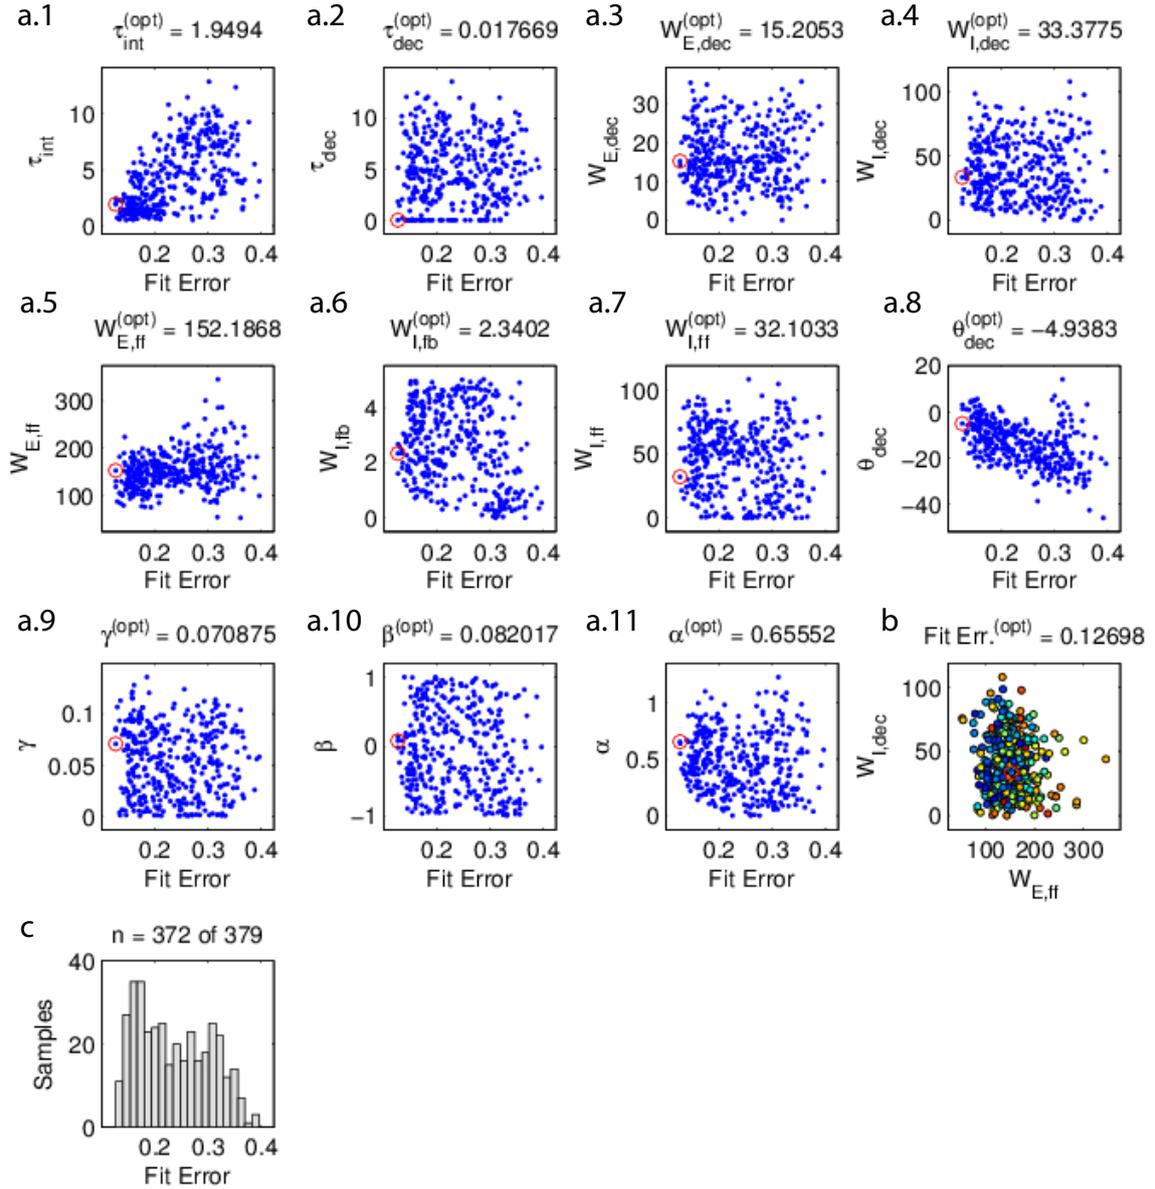


Figure B.1: Results of the optimization algorithm. **(a.1-11):** For each of the 11 parameter, blue dot corresponds to the value of the parameter at the outcome of one minimization, a red circle highlights to the optimal parameter value. **(a.1):** Evidence pools time-constant (s). **(a.2):** Decision pools time-constant (s). **(a.3):** Decision pools self-excitation. **(a.4):** Decision pools mutual-inhibition. **(a.5):** Specific feedforward excitation. **(a.6):** Specific feedback inhibition. **(a.7):** Unspecific feedforward inhibition. **(a.8):** Decision layer offset (constant input). **(a.9-11):** α , β and γ the input-mapping parameters. **(b)** Each colored dot corresponds to one minimization in the space $(W_{E,ff}, W_{I,dec})$, corresponding to the bifurcation parameters of the model. Colors represent the fit error (blue-to-red, low-to-high fit error). Minimizations form a dense group showing the oscillatory regime, where perceptual alternations are obtained. **(c):** Histogram of minimizations vs. fit error.

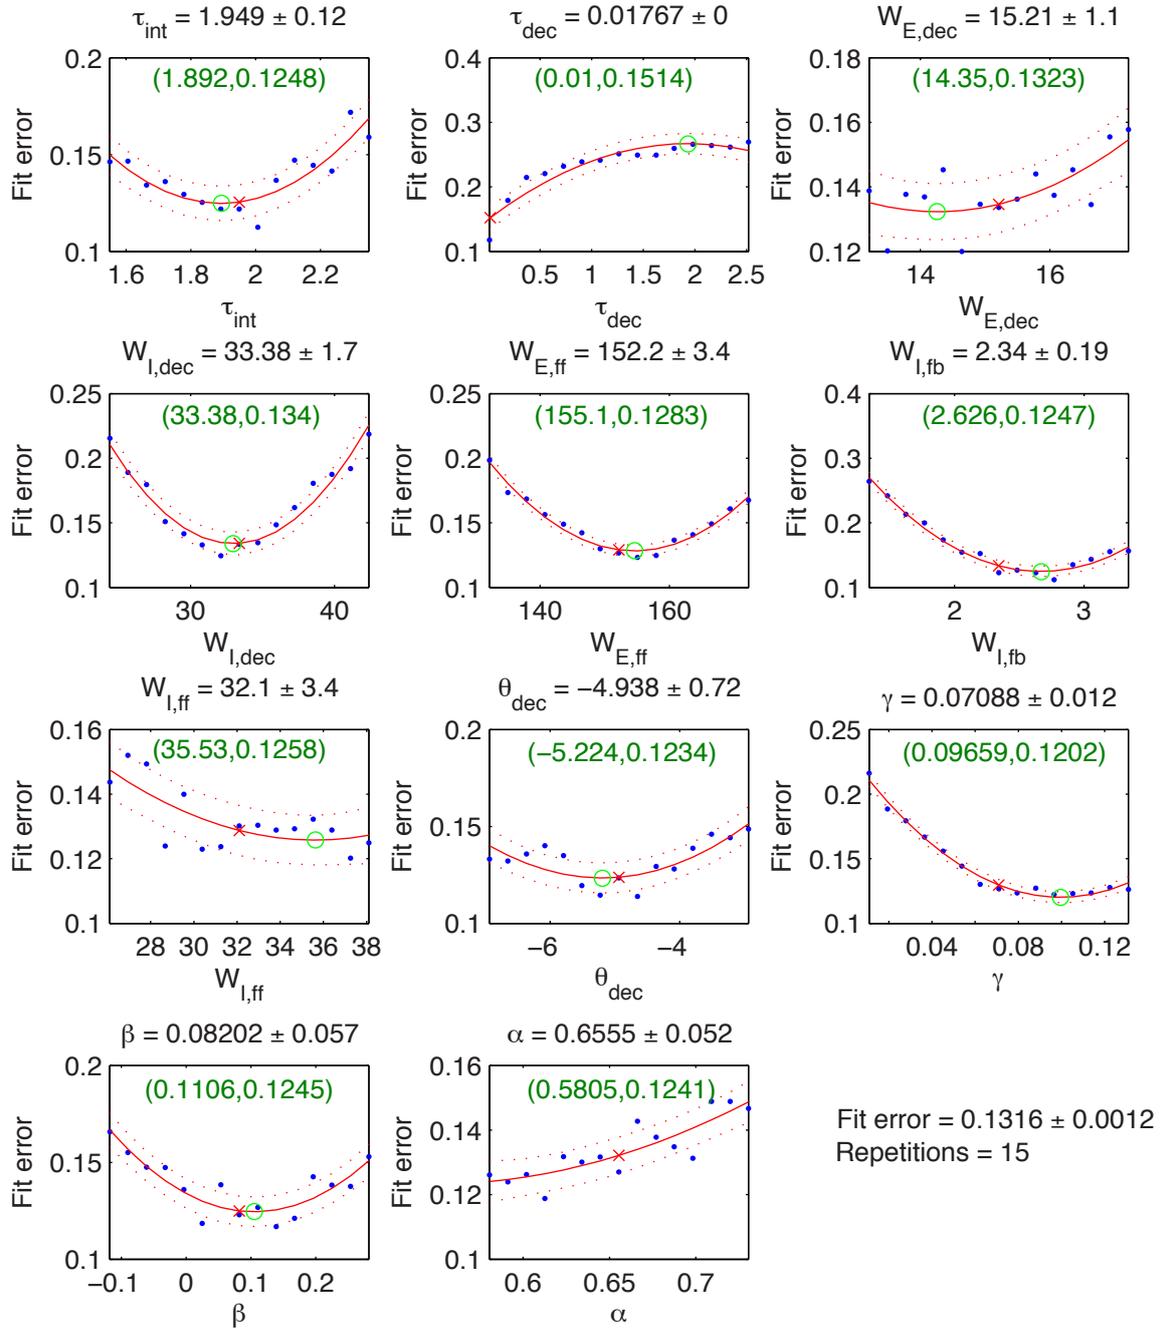


Figure B.2: For each panel, corresponding to each parameter α_i : **blue dots** correspond to the position of the pairs $(\alpha_i^{(j)}, \bar{R}^{(j)})$. A **red cross** indicates the position of the pair $(\alpha_i^{\text{opt}}, \bar{R}^{\text{opt}})$. **Red solid curves** represent the quadratic approximation of the scatter plot passing through $(\alpha_i^{\text{opt}}, \bar{R}^{\text{opt}})$, **red dotted curves** represent the standard deviation between the quadratic approximation and the $(\alpha_i^{(j)}, \bar{R}^{(j)})$ pairs. A **green circle** locates the position of the estimated extrema, whose is indicated in **green brackets** (left) with the value corresponding estimated fit error (right).

Transition-triggered averages

A convenient way to verify that the model operates in its intended regime, is to study transition-triggered averages, or in other words, the average values of activity in evidence and decision pools preceding and succeeding perceptual reversals, are . The following figures confirm that evidence pools cross the perceptual threshold in a DDR, driving a fast and sharp transition at the level of decision pools. The drift-dominated approach to the threshold is best appreciated when observing the dynamics of the evidence-bias $\Delta(t) = x(t) - y(t)$. The position of the differential threshold is well-defined and the reorganisation of sensory evidences, where the habituation and recovery phases are reversed, immediately follows perceptual reversals.

Three different situations are proposed: when both contrasts are low (see **Fig. C.1**), when both contrasts are high (see **Fig. C.2**), and when one contrast is low, and the other high (see **Fig. C.3**). Note that for increasing contrasts the average activity of evidence pools is increased, and reversal threshold decrease. Also, transitions are sharper for increasing contrast, as is made apparent by the slope of the evidence-bias at the threshold.

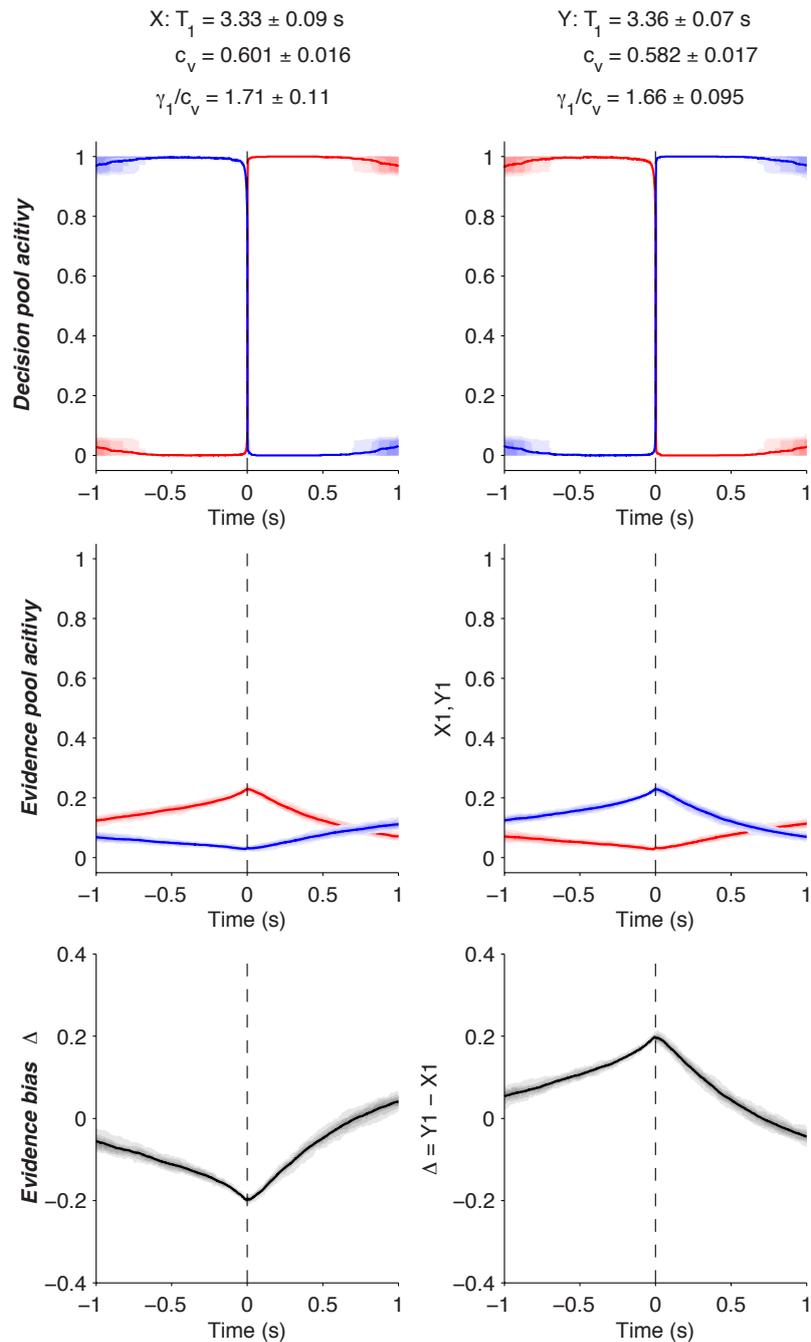


Figure C.1: Transition-triggered averages in the low contrasts condition (both contrasts set to 0.0625%), reversal times are centred at $t = 0$. Dominance duration statistics (mean, CV and skewness) for percept X and Y are indicated at the top. **Left column:** activities in pools associated with percept X (blue) losing dominance. **Right column:** activities in pools associated with percept Y (red) losing dominance. **Top:** Decision pools activity. **Middle:** Evidence pools activity. **Bottom:** Evidence-bias Δ .

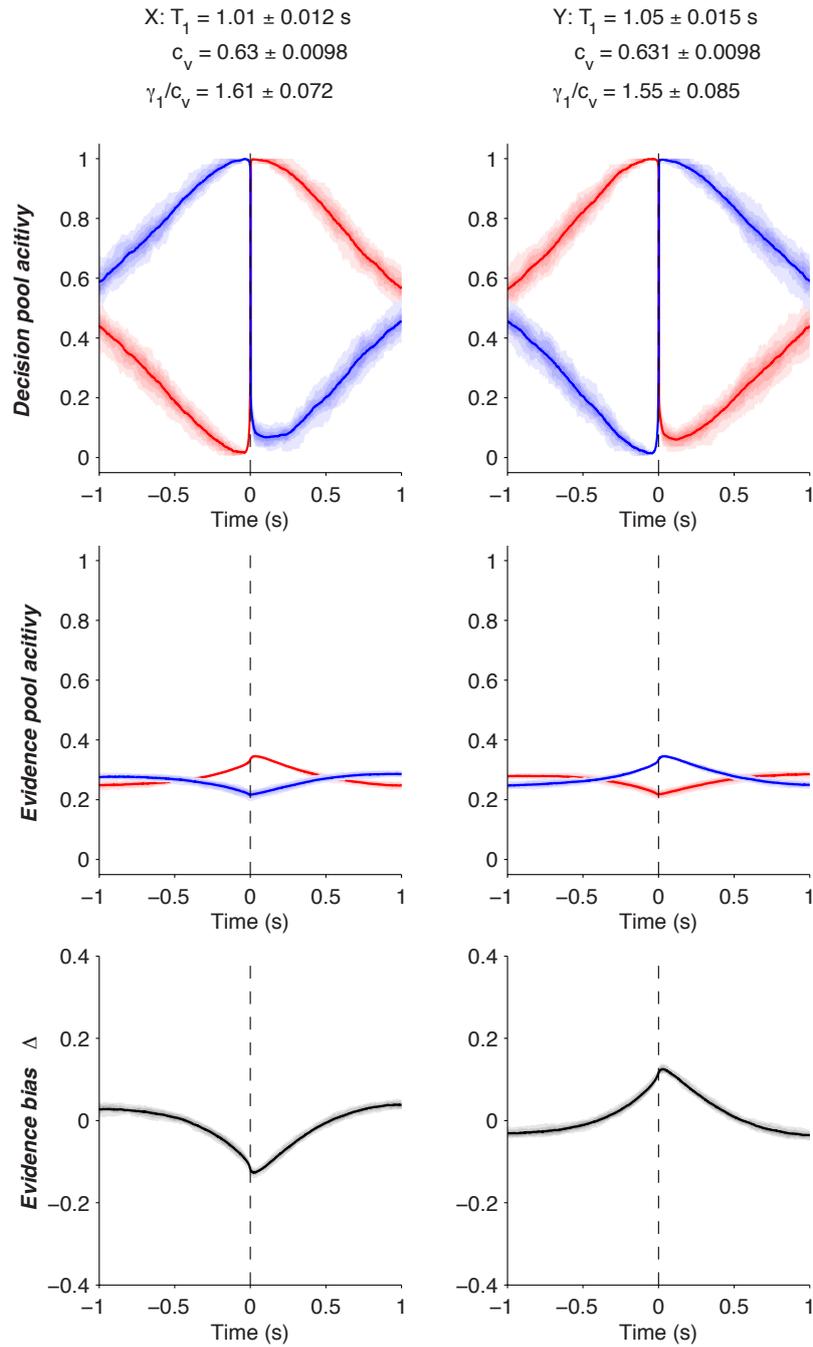


Figure C.2: Transition-triggered averages in the high contrasts condition (both contrasts set to 100%), reversal times are centred at $t = 0$. Dominance duration statistics (mean, CV and skewness) for percept X and Y are indicated at the top. **Left column:** activities in pools associated with percept X (blue) losing dominance. **Right column:** activities in pools associated with percept Y (red) losing dominance. **Top:** Decision pools activity. **Middle:** Evidence pools activity. **Bottom:** Evidence-bias Δ .

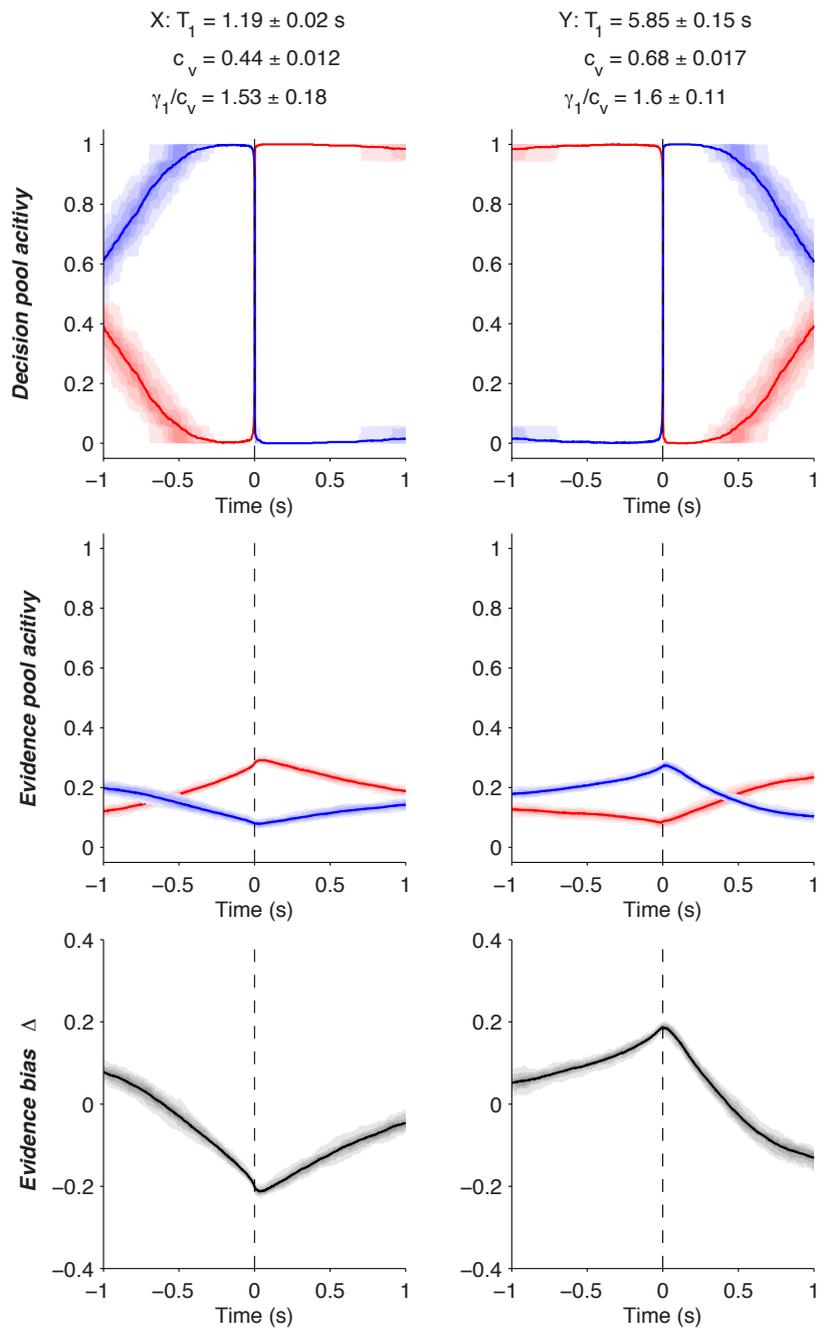


Figure C.3: Transition-triggered averages in the unbalanced contrasts condition (contrasts for percept X (blue) set to 0.0625%, contrast for percept Y (red) set to 100%), reversal times are centred at $t = 0$. Dominance duration statistics (mean, CV and skewness) for X and Y are indicated at the top, here X is weak, has short dominances and slightly lower CV; Y is strong, has long dominances and slightly greater CV. **Left column:** activities in pools associated with X losing dominance. **Right column:** activities in pools associated with Y losing dominance. **Top:** Decision pools activity. **Middle:** Evidence pools activity. **Bottom:** Evidence-bias Δ .

Bibliography

- M. Abeles, H. Bergman, I. Gat, I. Meilijson, E. Seidemann, N. Tishby, and E. Vaadia. Cortical activity flips among quasi-stationary states. *Proc. Natl. Acad. Sci. U.S.A.*, 92:8616–20, 1995.
- E. H. Adelson and J. A. Movshon. Phenomenal coherence of moving visual patterns. *Nature*, 300(5892):523–5, 1982.
- D. Amit. The Hebbian paradigm reintegrated: local reverberations as internal representations. *Behav. Brain Sci.*, 18:617–625, 1995.
- D. J. Amit and N. Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex*, 7:237–52, 1997.
- S. Amitay, J. Guiraud, E. Sohoglu, O. Zobay, B. A. Edmonds, Y. X. Zhang, and D. R. Moore. Human Decision Making Based on Variations in Internal Noise: An EEG Study. *PLoS ONE*, 8(7):1–14, 2013.
- A. Arieli, A. Sterkin, A. Grinvald, and A. Aertsen. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, 273(5283):1868–1871, Sep 1996.
- H. Barlow. Conditions for versatile learning, helmholtz’s unconscious inference, and the task of perception. *Vision Research*, 30(11):1561 – 1571, 1990.
- J. Beck, W. Ma, R. Kiani, T. Hanks, A. Churchland, J. Roitman, M. Shadlen, P. Latham, and A. Pouget. Probabilistic population codes for bayesian decision making. *Neuron*, 60:1142–52, 2008.
- J. M. Beggs and D. Plenz. Neuronal avalanches in neocortical circuits. *J. Neurosci.*, 23(35):11167–77, 2003.
- J. M. Beggs and D. Plenz. Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *J. Neurosci.*, 24(22):5216–29, 2004.
- P. Berkes, G. Orbán, M. Lengyel, and J. Fiser. Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment. *Science*, 331(January):83–88, 2011.
- N. H. Bingham. Fluctuation theory for the ehrenfest urn. *Adv. Appl. Probab.*, 23(3):598–611, 1991.

- S. Bitzer, H. Park, F. Blankenburg, and S. J. Kiebel. Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. *Frontiers in human neuroscience*, 8:102, 2014.
- R. Blake and N. K. Logothetis. Visual competition. *Nat. Rev. Neurosci.*, 3:13–21, 2002.
- R. R. Blake, R. Fox, and C. McIntyre. Stochastic properties of stabilized-image binocular rivalry alternations. *J. Exp. Psychol.*, 88(3):327–332., 1971.
- R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700–765, 2006.
- A. Bollimunta and J. Ditterich. Local computation of decision-relevant net sensory evidence in parietal cortex. *Cerebral Cortex*, 22(4), 2012.
- Y. S. Bonnef, A. Cooperman, and D. Sagi. Motion-induced blindness in normal observers. *Nature*, 411(6839):798–801, 2001.
- A. Borsellino, A. De Marco, A. Allazetta, S. Rinesi, and B. Bartolini. Reversal time distribution in the perception of visual ambiguous stimuli. *Kybernetik*, 10:139–44, 1972.
- J. Brascamp, P. Klink, and W. Levelt. The laws of binocular rivalry: 50 years of levelt’s propositions. *Vision Research*, 109, Part A:20 – 37, 2015.
- J. W. Brascamp, R. van Ee, W. R. Pestman, and A. V. van den Berg. Distributions of alternation rates in various forms of bistable perception. *J. Vision*, 5(4):287–98, 2005.
- J. W. Brascamp, R. van Ee, A. J. Noest, R. H. A. H. Jacobs, and A. V. Van den Berg. The time course of binocular rivalry reveals a fundamental role of noise. *J. Vision*, 6(11):1244–1256, 2006.
- J. Braun and M. Mattia. Attractors and noise: Twin drivers of decisions and multistability. *NeuroImage*, 52(3):740–751, 2010.
- A. S. Bregman. *Auditory scene analysis: The perceptual organization of sound*. MIT press, 1994.
- N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons.. *J. Comput. Neurosci.*, 8:183–208, 2000.
- N. Brunel and V. Hakim. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.*, 11:1621–71, 1999.
- M. A. Buice and C. C. Chow. Dynamic Finite Size Effects in Spiking Neural Networks. *PLoS Computational Biology*, 9(1), 2013.
- F. W. Campbell and E. R. Howell. Monocular alternation: a method for the investigation of pattern vision. *J. Physiol.*, 225(2):19P–21P, 1972.
- R. Cao, J. Braun, and M. Mattia. Stochastic accumulation by cortical columns may explain the scalar property of multi-stable perception. *Phys. Rev. Lett.*, 113(9):098103, 2014.

- R. Cao, A. Pastukhov, M. Mattia, and J. Braun. Reverse-engineering perceptual inference as a hierarchical interaction between spontaneously active cortical columns. In *Proceedings of the Bernstein Conference*, 2015a. doi: 10.12751/nmcn.bc2015.0018.
- R. Cao, A. Pastukhov, M. Mattia, and J. Braun. An experimentally constrained theory for Levels propositions and the Scalar Property of multistable perception. In *Perception*, volume 44, page 226, 2015b.
- R. Cao, A. Pastukhov, M. Mattia, and J. Braun. Collective Activity of Many Bistable Assemblies Reproduces Characteristic Dynamics of Multistable Perception. *The Journal of Neuroscience*, 36(26):6957–6972, 2016.
- R. H. S. Carpenter. Analysing the detail of saccadic reaction time distribution. *Biocybern. Biomed. Eng.*, 32(2):49–63, 2012.
- O. Carter, T. Konkle, Q. Wang, V. Hayward, and C. Moore. Tactile rivalry demonstrated with an ambiguous apparent-motion quartet. *Curr. Biol.*, 18(14):1050–4, 2008.
- A. K. Churchland, R. Kiani, R. Chaudhuri, X. J. Wang, A. Pouget, and M. N. Shadlen. Variance as a signature of neural computations during decision making. *Neuron*, 69(4):818–831, 2011.
- D. R. Cox and H. D. Miller. *The Theory of Stochastic Processes*. Chapman and Hall, London, 1972.
- J. C. Cox, J. E. Ingersoll, and S. A. Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407, 1985.
- R. Curtu, A. Shpiro, N. Rubin, and J. Rinzel. Mechanisms for Frequency Control in Neuronal Competition Models. *SIAM J. Appl. Dynam. Sys.*, 7(2):609–649, 2008.
- A. De Marco, P. Penengo, A. Trabucco, A. Borsellino, F. Carlini, M. Piani, and M. T. Tuccio. Stochastic models and fluctuations in reversal time of ambiguous figures. *Perception*, 6(6): 645–656, 1977.
- G. Deco, M. Perez-Sanagustin, V. de Lafuente, and R. Romo. Perceptual detection as a dynamical bistability phenomenon: a neurocomputational correlate of sensation. *Proc. Natl. Acad. Sci. U.S.A.*, 104(50):20073–20077, 2007.
- M. Deger, T. Schwalger, R. Naud, and W. Gerstner. Fluctuations and information filtering in coupled populations of spiking neurons with adaptation. *Phys. Rev. E*, 90, 2014. doi: 10.1103/PhysRevE.90.062704.
- S. Dehaene and J. P. Changeux. Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, 70(2):200–227, 2011.
- S. Deneve. Bayesian spiking neurons I: inference. *Neural computation*, 20(1):91–117, 2008a.
- S. Deneve. Bayesian spiking neurons II: learning. *Neural Computation*, 20(1):118–145, 2008b.
- S. Deneve. Making decisions with unknown sensory reliability. *Frontiers in Neuroscience*, 6:1–13, 2012.

- A. Destexhe and D. Paré. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J. Neurophysiol.*, 81(4):1531–1547, 1999.
- H. Dette. On a generalization of the Ehrenfest urn model. *J. Appl. Probab.*, 31(4):930–939, 1994.
- J. Ditterich. A comparison between mechanisms of multi-alternative perceptual decision making: Ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Frontiers in Neuroscience*, 4:1–24, 2010.
- R. Douglas, C. Koch, M. Mahowald, K. Martin, and H. Suarez. Recurrent excitation in neocortical circuits. *Science*, 269(5226):981–985, 1995.
- R. J. Douglas and K. A. C. Martin. Recurrent neuronal circuits in the neocortex. *Current Biology*, 17(13):496–500, 2007.
- J. Drugowitsch, R. Moreno-Bote, A. K. Churchland, M. N. Shadlen, and A. Pouget. The cost of accumulating evidence in perceptual decision making. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32(11):3612–3628, 2012.
- D. Durstewitz and G. Deco. Computational significance of transient dynamics in cortical networks. *Eur. J. Neurosci.*, 27:217–27, 2008.
- P. Ehrenfest. Über zwei bekannte Einwände gegen das Boltzmannsche H-Theorem. *Physikalische Zeitschrift*, 8:311–4, 1907.
- A. A. Faisal, L. P. J. Selen, and Wolpert. Noise in the nervous system. *Nat. Rev. Neurosci.*, 9(4):292–303, 2008.
- J. Fiser, P. Berkes, G. Orban, and M. Lengyel. Statistically optimal perception and learning : from behavior to neural representations. *Trends in Cognitive Sciences*, pages 119–130, 2010.
- M. B. Flegg, P. K. Pollett, and D. K. Gramotnev. Ehrenfest model for condensation and evaporation processes in degrading aggregates with multiple bonds. *Phys. Rev. E*, 78(3 Pt 1):031117, 2008.
- R. Fox and J. Herrmann. Stochastic properties of binocular rivalry alternations. *Percept. Psychophys.*, 2:432–6, 1967.
- F. Freyer, K. Aquino, P. A. Robinson, P. Ritter, and M. Breakspear. Bistability and Non-Gaussian Fluctuations in Spontaneous Cortical Activity. *J. Neurosci.*, 29:8512?–24, 2009.
- F. Freyer, J. A. Roberts, P. Ritter, and M. Breakspear. A canonical model of multistability and scale-invariance in biological systems. *PLoS Comput Biol*, 8(8):1–15, 08 2012. doi: 10.1371/journal.pcbi.1002634.
- K. Friston. Learning and inference in the brain. *Neural Networks*, 16(9):1325–1352, 2003.
- K. Friston. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.*, 11(2):127–38, 2010.
- K. Friston, M. Breakspear, and G. Deco. Perception and self-organized instability. *Frontiers in Computational Neuroscience*, 6(July):1–19, 2012. ISSN 1662-5188.

- C. W. Gardiner. *Handbook of stochastic methods for physics, chemistry and the natural sciences*. Springer-Verlag, New York, 1985.
- S. Gershman, E. Vul, and J. B. Tenenbaum. Perceptual multistability as Markov chain Monte Carlo inference. *Nips*, pages 1–9, 2009.
- S. J. Gershman, E. Vul, and J. B. Tenenbaum. Multistability and perceptual inference. *Neural Comput.*, 24(1):1–24, 2012.
- K. Ghosh, K. A. Dill, M. M. Inamdar, E. Seitaridou, and R. Phillips. Teaching the principles of statistical dynamics. *Am. J. Phys.*, 74:123–33, 2006.
- G. Gigante, M. Mattia, J. Braun, and P. Del Giudice. Bistable perception modeled as competing stochastic integrations at two levels. *PLoS Comput. Biol.*, 5:e1000430, 2009.
- J. Gold and M. Shadlen. The neural basis of decision making. *Annu. Rev. Neurosci.*, 30:535–74, 2007.
- I. S. Gradshteyn and I. M. Ryzhik. *Tables of Integrals, Series, and Products*, volume 56. 1988.
- D. M. Green and J. a. Swets. *Signal detection theory and psychophysics*. New York: John Wiley & Sons, Inc., 1966.
- R. Gregory. *Seeing Through Illusions*. Oxford University Press, USA, 2009.
- R. M. Haefner, P. Berkes, and J. Fiser. Perceptual Decision-Making as Probabilistic Inference by Neural Sampling. *Neuron*, 90(3):649–660, 2016.
- F. Han, N. Caporale, and Y. Dan. Reverberation of Recent Visual Experience in Spontaneous Cortical Waves. *Neuron*, 60(2):321–327, 2008.
- P. Hanggi, Peter, and H. Thomas. Bistable systems: Master equation versus Fokker-Planck modeling. *Phys. Rev. A*, 29(1):371–378, 1984.
- P. Hanggi, P. Talkner, and M. Borkovec. Reaction-rate theory: fifty years after Kramers. *Rev. Mod. Phys.*, 62(2):251–341, 1990.
- T. D. Hanks, M. E. Mazurek, R. Kiani, E. Hopp, and M. N. Shadlen. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(17):6339–6352, 2011.
- J. Hesse and T. Gross. Self-organized criticality as a fundamental property of neural systems. *Frontiers in Systems Neuroscience*, 8(September):166, 2014.
- G. Hinton. Where do features come from? *Cognitive Science*, 38(6):1078–1101, 2014.
- G. E. Hinton and T. J. Sejnowski. Learning and relearning in Boltzmann machines. In G. Rumelhart and J. McClelland, editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations*. MIT Press, 1986.
- J. Hohwy, A. Roepstorff, and K. Friston. Predictive coding explains binocular rivalry: an epistemological review. *Cognition*, 108(3):687–701, 2008.

- M. Hollins. The effect of contrast on the completeness of binocular rivalry suppression. *Percept. & Psychophys.*, 27(6):550–6, 1980.
- P. O. Hoyer and A. Hyvärinen. Interpreting neural response variability as monte carlo sampling of the posterior. *Adv. Neur. Inf. Proc. Sys.*, 15:293–300, 2003.
- J. Inoue, S. Sato, and L. Ricciardi. On the parameter estimation for diffusion models of single neuron’s activities. *Biol. Cybern.*, 73(3):209–221, 1995.
- E. T. Jaynes. *E.T. Jaynes : papers on probability, statistics, and statistical physics*. Synthese library. D. Reidel Hingham, MA, 1983. ISBN 90-277-1448-7.
- R. Kanai, D. Carmel, B. Bahrami, and G. Rees. Structural and functional fractionation of right superior parietal cortex in bistable perception. *Curr. Biol.*, 21(3):R106–R107, 2011.
- M.-S. Kang. Size matters: a study of binocular rivalry dynamics. *Journal of vision*, 9(1):17.1–11, 2009.
- M.-S. Kang and R. Blake. What causes alternations in dominance during binocular rivalry? *Atten. Percept. & Psychophys.*, 72(1):179–86, 2010.
- S. Karlin and J. L. McGregor. A characterization of birth and death processes. *Proc. Natl. Acad. Sci. USA*, 45(3):375–379, 1959.
- S. Karlin and J. L. McGregor. Ehrenfest Urn Models. *J. Appl. Probab.*, 2(2):352–376, 1965.
- T. Kenet, D. Bibitchkov, M. Tsodyks, A. Grinvald, and A. Arieli. Spontaneously emerging cortical representations of visual attributes. *Nature*, 425:954–6, 2003.
- D. Kersten, P. Mamassian, and A. Yuille. Object perception as bayesian inference. *Annu. Rev. Psychol.*, 55:271–304, 2004.
- J. Kim and M. Shadlen. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.*, 2:176–85, 1999.
- Y.-J. Kim, M. Grabowecky, and S. Suzuki. Stochastic resonance in binocular rivalry. *Vision Res.*, 46:392–406, 2006.
- A. Kleinschmidt, P. Sterzer, and G. Rees. Variability of perceptual multistability: from brain state to individual trait. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367:988–1000, 2012.
- P. C. Klink, R. van Ee, and R. J. a. van Wezel. General validity of levelt’s proportions reveals common computational mechanisms for visual rivalry. *PLoS ONE*, 3(10):1–9, 2008.
- T. H. J. Knapen, J. W. Brascamp, J. Pearson, R. van Ee, and R. Blake. The Role of Frontal and Parietal Brain Areas in Bistable Perception. *J. Neurosci.*, 31(28):10293–10301, 2011.
- D. C. Knill and A. Pouget. The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, 2004.

- J. Kornmeier and M. Bach. Ambiguous figures - what happens in the brain when perception changes but not the stimulus. *Frontiers in human neuroscience*, 6:51, 2012.
- J. Kornmeier, C. M. Hein, and M. Bach. Multistable perception: when bottom-up and top-down coincide. *Brain Cogn.*, 69(1):138–47, 2009.
- H. Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions., 1940. ISSN 00318914.
- C. R. Laing and C. C. Chow. A spiking neuron model for binocular rivalry. *J. Comput. Neurosci.*, 12:39–53, 2002.
- P. Lánský and L. Sacerdote. The OrnsteinUhlenbeck neuronal model with signal-dependent noise. *Physics Letters A*, 285:132140, 2001.
- T. Larsen and R. Bogacz. Initiation and termination of integration in a decision process. *Neural Networks*, 23(3):322–333, 2010.
- K. W. Latimer, J. L. Yates, M. L. R. Meister, A. C. Huk, and J. W. Pillow. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244):184–7, 2015.
- S. R. Lehky. An astable multivibrator model of binocular rivalry. *Perception*, 17:215–28, 1988.
- S. R. Lehky. Binocular rivalry is not chaotic. *Proc. Roy. Soc. Lond. B*, 259:71–6, 1995.
- D. A. Leopold and N. K. Logothetis. Multistable phenomena: changing views in perception. *Trends Cogn. Sci.*, 3(7):254–264, 1999.
- D. A. Leopold and A. Maier. Ongoing physiological processes in the cerebral cortex. *Neuroimage*, 62(4):2190–2200, 2012.
- W. J. M. Levelt. *On binocular rivalry*. Van Gorkum Comp., Leiden, 1965.
- W. J. M. Levelt. Note on the distributions of dominance times in binocular rivalry. *Br. J. Psychol.*, 58:143–5, 1967.
- A. Litwin-Kumar and B. Doiron. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.*, 15(11):1498–1505, 2012.
- A. Luczak, P. Barth, and H. K. D. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–425, 2009.
- E. D. Lumer and G. Rees. Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc. Natl. Acad. Sci. U.S.A.*, 96(4):1669–73, 1999.
- J. S. Lund, A. Angelucci, and P. C. Bressloff. Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cereb. Cortex*, 13(1):15–24, 2003.
- W. Ma, J. Beck, P. Latham, and A. Pouget. Bayesian inference with probabilistic population codes. *Nat. Neurosci.*, 9:1432–8, 2006.

- G. Major and D. Tank. Persistent neural activity: prevalence and mechanisms. *Curr. Opin. Neurobiol.*, 14(6):675–684, Dec 2004.
- M. Mattia and P. Del Giudice. Population dynamics of interacting spiking neurons. *Phys. Rev. E*, 66:051917, 2002.
- M. Mattia, P. Pani, G. Mirabella, S. Costa, P. Del Giudice, and S. Ferraina. Heterogeneous attractor cell assemblies for motor planning in premotor cortex. *J. Neurosci.*, 33(27):11155–68, 2013.
- M. E. Mazurek, J. D. Roitman, J. Ditterich, and M. N. Shadlen. A Role for Neural Integrators in Perceptual Decision Making. *Cerebral Cortex*, 13(11):1257–1269, 2003.
- F. Megumi, B. Bahrami, R. Kanai, and G. Rees. Brain activity dynamics in human parietal regions during spontaneous switches in bistable perception. *NeuroImage*, 107:190–197, 2015.
- M. Meng and F. Tong. Can attention selectively bias bistable perception? Differences between binocular rivalry and ambiguous figures. *J. Vision*, 4(7):539–51, 2004.
- R. Moran. Optimal decision making in heterogeneous and biased environments. *Psychonomic bulletin & review*, 22(1):38–53, 2015.
- R. Moreno-Bote and A. Knill, Pouget. Bayesian sampling in visual perception. *Proc. Nat. Acad. Sci. U.S.A.*, pages 1–6, 2011.
- R. Moreno-Bote, J. Rinzel, and N. Rubin. Noise-induced alternations in an attractor network model of perceptual bistability. *J. Neurophysiol.*, 98:1125–39, 2007.
- R. Moreno-Bote, A. Shpiro, J. Rinzel, and N. Rubin. Alternation rate in perceptual bistability is maximal at and symmetric around equi-dominance. *J. Vision*, 10(11):1, 2010.
- T. Murata, N. Matsui, S. Miyauchi, Y. Kakita, and T. Yanagida. Discrete stochastic process underlying perceptual rivalry. *Neuroreport*, 14(10):1347–52, 2003.
- M. Nawrot and R. Blake. Neural integration of information specifying structure from stereopsis and motion. *Science*, 244(4905):716–8, 1989.
- H. Nienborg and B. G. Cumming. Decision-related activity in sensory neurons reflects more than a neuron’s causal effect. *Nature*, 459(7243):89–92, 2009.
- H. Okamoto and T. Fukai. Neural mechanism for a cognitive timer. *Phys. Rev. Lett.*, 86(17):3919–22, 2001.
- G. Orbán, P. Berkes, J. Fiser, and M. Lengyel. Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex. *Neuron*, 92(2):530–543, 2016.
- J. L. Palacios. Fluctuation theory for the ehrenfest urn via electric networks. *Adv. Appl. Probab.*, 25:472–476, 1993.
- A. Pastukhov and J. Braun. Perceptual reversals need no prompting by attention. *J. Vision*, 7(10):5.1–17, 2007.

- A. Pastukhov and J. Braun. Cumulative history quantifies the role of neural adaptation in multistable perception. *J. Vision*, 11(10):12, 2011.
- A. Pastukhov and J.-N. Klanke. Exogenously triggered perceptual switches in multistable structure-from-motion occur in the absence of visual awareness. *Journal of vision*, 16(3):14, 2016.
- A. Pastukhov, P. E. García-Rodríguez, J. Haenicke, A. Guillamon, G. Deco, and J. Braun. Multistable perception balances stability and sensitivity. *Frontiers in Computational Neuroscience*, 7(March):17, 2013. ISSN 1662-5188.
- B. Pearson, J. Raskevicius, P. M. Bays, Y. Pertzov, and M. Husain. Working memory retrieval as a decision process. *J. Vision*, 14(2):1–15, 2014.
- T. Petermann, T. C. Thiagarajan, M. A. Lebedev, M. A. Nicolelis, D. R. Chialvo, and D. Plenz. Spontaneous cortical activity in awake monkeys composed of neuronal avalanches. *Proc. Natl. Acad. Sci. U.S.A.*, 106(37):15921–6, 2009.
- T. J. Petersik. Buildup and decay of a three-dimensional rotational aftereffect obtained with a three-dimensional figure. *Perception*, 31(7):825–36, 2002.
- A. Pouget, J. M. Beck, W. J. Ma, and P. E. Latham. Probabilistic brains: knowns and unknowns. *Nature neuroscience*, 16(9):1170–8, 2013.
- S. Pressé, K. Gosh, J. Lee, and K. A. Dill. Principles of maximum entropy and maximum caliber in statistical physics. *Reviews of Modern Physics*, 85:1115–41, 2013.
- D. Pressnitzer and J. M. Hupe. Temporal dynamics of auditory and visual bistability reveal common principles of perceptual organization. *Curr. Biol.*, 16(13):1351–7, 2006.
- K. H. Pribram. *Origins: Brain and Self-Organization*. Lawrence Erlbaum, 1994.
- R. Ratcliff and G. McKoon. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20:873–922, 2008.
- R. Ratcliff and P. L. Smith. A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.*, 111(2):333–367, 2004.
- H. Risken. Fokker-planck equation. In *The Fokker-Planck Equation*, volume 18 of *Springer Series in Synergetics*, pages 63–95. Springer Berlin Heidelberg, 1984.
- C. P. Robert and G. Casella. Monte carlo statistical methods, 1998.
- J. Roitman and M. Shadlen. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.*, 22:9475–89, 2002.
- T. L. Saaty. Some stochastic processes with absorbing barriers. *J. R. Stat. Soc. Ser. B*, 23(2): 319–334, 1961.
- L. Sacerdote and P. Lansky. Interspike inter val statistics in the Ornstein-Uhlenbeck neuronal model with signal-dependent noise. *BioSystems*, 67:213–219, 2002.

- J. D. Schall. Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.*, 2(1):33–42, 2001.
- T. Schwalger. The Interspike-Interval Statistics of Non-Renewal Neuron Models. *PhD dissertation*, 2013.
- J. Seely and C. C. Chow. Role of mutual inhibition in binocular rivalry. *J. Neurophysiol.*, 106(5): 2136–50, 2011.
- M. N. Shadlen and R. Kiani. Decision making as a window to cognition. *Neuron*, 80(3):791–806, 2013.
- A. Shpiro, R. Curtu, J. Rinzel, and N. Rubin. Dynamical characteristics common to neuronal competition models. *J. Neurophysiol.*, 97(1):462–73, 2007.
- A. Shpiro, R. Moreno-Bote, N. Rubin, and J. Rinzel. Balance between noise and adaptation in competition models of perceptual bistability. *J. Comput. Neurosci.*, pages 37–54, 2009.
- P. L. Smith. Stochastic dynamic models of response time and accuracy: A foundational primer. *J. Math. Psychol.*, 44(3):408–463, 2000.
- P. L. Smith and R. Ratcliff. An integrated theory of attention and decision making in visual signal detection. *Psychol. Rev.*, 116:283–317, 2009.
- P. L. Smith, R. Ratcliff, and G. McKoon. The diffusion model is not a deterministic growth model: comment on jones and dzhafarov (2014). *Psychol. Rev.*, 121(4):679–688, 2014.
- J. C. Spall. *Introduction to Stochastic Search and Optimization*. John Wiley & Sons, Inc., New York, NY, USA, 1 edition, 2003. ISBN 0471330523.
- G. Sperling and B. A. Doshier. Depth from motion. In T. V. Papathomas, A. G. Charles Chubb, and E. Kowler, editors, *Early Vision and Beyond*, pages 133–142. MIT Press, Cambridge, MA, 1994.
- P. Sterzer and G. Rees. A neural basis for percept stabilization in binocular rivalry. *J. Cogn. Neurosci.*, 20(3):389–99, 2008.
- P. Sterzer, A. Kleinschmidt, and G. Rees. The neural bases of multistable perception. *Trends Cogn. Sci.*, 13(7):310–8, 2009.
- L. P. Sugrue and W. T. Corrado. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.*, 6(5):363–375, 2005.
- R. Sundaeswara and P. R. Schrater. Perceptual multistability predicted by search model for Bayesian decisions. *J. Vision*, 8(5):1–19, 2008.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The Mit Press, 1998.
- J. A. Swets. *Signal Detection and Recognition in Human Observers: Contemporary Readings*. John Wiley and Sons, 1964.
- H. Tanigawa, Q. Wang, and I. Fujita. Organization of horizontal axons in the inferior temporal cortex and primary visual cortex of the macaque monkey. *Cereb. Cortex*, 15(12):1887–1899, 2005.

- M. Taylor and K. Aldridge. Stochastic processes in reversing figure perception. *Percept. & Psychophys.*, 16(1):9–25, 1974.
- M. Tsodyks, T. Kenet, and A. Arieli. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, 286(5446):1943–6, 1999.
- H. C. Tuckwell. *Introduction to theoretical neurobiology: Volume 2, Nonlinear and stochastic theories*. Cambridge Univ. Press, Cambridge, 1988/2008.
- R. van Ee. Stochastic variations in sensory awareness are driven by noisy neuronal adaptation: evidence from serial correlations in perceptual bistability. *J. Opt. Soc. Am. A*, 26(12):2612–22, 2009.
- N. G. van Kampen. *Stochastic processes in physics and chemistry*. North-Holland Physics Publishing, Amsterdam, 1981.
- A. M. Van Loon, T. Knapen, H. S. Scholte, E. St. John-Saaltink, T. H. Donner, and V. A. F. Lamme. GABA shapes the dynamics of bistable perception. *Current Biology*, 23(9):823–827, 2013.
- E. Vargo and L. Leemis. Moment-Ratio Diagrams for Univariate Distributions. *Journal of Quality Technology*, 42(3):276–286, 2010.
- A. Veliz-Cuba, Z. P. Kilpatrick, and K. Josić. Stochastic Models of Evidence Accumulation in Changing Environments. *SIAM Review*, 58(2):264–289, 2016.
- M. von Gruenau and S. Dube. Ambiguous plaids: switching between coherence and transparency. *Spatial Vision*, 7(3):199–211, 2005.
- H. von Helmholtz. *Treatise on Physiological Optics: Vol. 3.*, volume 3. The Optical Society of America, Birmingham, Alabama, 1866.
- A. Wald. *Sequential Analysis*. John Wiley and Sons, 1st edition, 1947.
- P. Walker. Stochastic properties of binocular rivalry alternations. *Percept. & Psychophys.*, 18(6):467–473, 1975.
- H. Wallach and D. N. O’Connell. The kinetic depth effect. *J. Exp. Psychol.*, 45(4):205–17, 1953.
- M. G. Wang, D. Arteaga, and B. J. J. He. Brain mechanisms for simple perception and bistable perception. *Proc Natl Acad Sci U S A*, 110, 2013.
- X.-J. Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36:955–68, 2002.
- X.-J. Wang. Neural dynamics and circuit mechanisms of decision-making. *Curr. Opin. Neurobiol.*, 22(6):1039–1046, 2012.
- T. Watanabe, N. Masuda, F. Megumi, R. Kanai, and G. Rees. Energy landscape and dynamics of brain activity during human bistable perception. *Nat. Commun.*, 5:1–11, 2014.

- J. A. White, J. T. Rubinstein, and A. R. Kay. Channel noise in neurons. *Trends in Neurosciences*, 23(3):131–137, 2000.
- H. R. Wilson. Computational evidence for a rivalry hierarchy in vision. *Proc. Natl. Acad. Sci. U.S.A.*, 100:14499–503, 2003.
- H. R. Wilson. Minimal physiological conditions for binocular rivalry and rivalry memory. *Vision Res.*, 47:2741–50, 2007.
- K. Wimmer, A. Compte, A. Roxin, D. Peixoto, A. Renart, and J. de la Rocha. Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nature communications*, 6:6177, 2015.
- I. Winkler, S. Denham, R. Mill, T. Bohm, and A. Bendixen. Multistability in auditory stream segregation: a predictive coding view. *Philos. Trans. Roy. Soc. Lond. B*, 367(1591):1001–12, 2012.
- J. M. Wolfe. Reversing ocular dominance and suppression in a single flash. *Vision Res.*, 24(5):471–8, 1984.
- C. Xenophontos. A formula for the n^{th} derivative of the quotient of two functions. Private communication, 2007.
- A. Yuille and D. Kersten. Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, 10(7):301–308, 2006.
- Y. H. Zhou, J. B. Gao, K. D. White, I. Merk, and K. Yao. Perceptual dominance time distributions in multistable visual perception. *Biol. Cybern.*, 90(4):256–263, 2004.