

How Do Computational Models of the Role of Dopamine as a Reward Prediction Error Map on to Current Dopamine Theories of Schizophrenia?

Angela J. Thurnham* (a.j.thurnham@herts.ac.uk), D. John Done** (d.j.done@herts.ac.uk),

Neil Davey* (n.davey@herts.ac.uk), Ray J. Frank* (r.j.frank@herts.ac.uk)

School of Computer Science,* School of Psychology, **University of Hertfordshire,
College Lane, Hatfield, Hertfordshire. AL10 9AB United Kingdom

Abstract

A review of the current dopamine theories of schizophrenia reveals a likely imbalance between cortical and subcortical microcircuits due to an insufficient inhibitory brake, leading to a disruption of the dopamine system and the classic positive psychotic symptoms, negative symptoms and cognitive deficits associated with the disorder. Recent computational models have modelled the role of dopamine as a reward prediction error, using Temporal Difference and have successfully shown how these symptoms could arise from a disturbance to the dopamine system. We review these models in the light of dopamine theories of schizophrenia and highlight some of the major points that should be addressed by future computational models.

Keywords: Dopamine; Schizophrenia; Neurocomputational Modelling; Saliency; Temporal Difference.

Theories of the role of dopamine over the last five years tend to converge on the idea that dopamine encodes a reward prediction error (RPE) of the discrepancy between actual and expected future reward. This discrepancy is used to drive learning towards actions which are necessary for survival in the real world (Schultz, 1998), and it is likely that disruption to this system gives rise to an abnormality in information processing by dopamine and some of the symptoms currently associated with schizophrenia, particularly psychosis and deficits in working memory. Temporal Difference Learning (Sutton, 1988; Sutton & Barto, 1998), a form of Reinforcement Learning Theory, provides an explicit method of modelling and quantifying the Reward Prediction, or Temporal Difference (TD), error (Schultz, Dayan & Montague, 1997; Hollerman & Schultz, 1998) and can be used as a valid computational implementation of the RPE for neural network simulations. While dopamine should not be viewed in isolation, but seen to be working in concert with other neurotransmitters, such as glutamate and GABA (Abi-Dhargham, 2004; Carlsson, Waters, Holm-Waters, Tedroff, Nilsson & Carlsson, 2001; Winterer & Weinberger, 2004), there are still attributes and deficiencies that can be strongly linked to dopamine activity.

The role of dopamine, and the possible location and nature of the dysfunction, presented in theories of schizophrenia by Carlsson et al, (2001); Kapur, (2003); Abi-Dhargham, (2004) and Winterer & Weinberger, (2004), are discussed in the first section on dopamine theories of schizophrenia below. The second section relates specifically to computational models, particularly existing connectionist models of dopamine as a

reward prediction, or TD error, including evidence that the RPE model of dopamine activity applies to humans as well as primates. The biological plausibility of existing neural network models by Cohen & Servan-Schreiber, (1992); Braver Barch & Cohen, (1999); Suri & Schultz, (1999); Rougier, Noelle, Braver, Cohen & O'Reilly, (2005) and O'Reilly & Frank, (2006) are then discussed in the light of the afore-mentioned dopamine theories of schizophrenia. Finally, we conclude with four major questions arising from recent dopamine theories of schizophrenia that remain to be addressed by current computational models.

Dopamine Theories of Schizophrenia

Role of Dopamine

It is generally agreed that dopamine enables the ability to focus on task relevant information. Current theories of the effects of dopamine on behaviour focus on the role of dopamine as a neuromodulator in Reinforcement Learning, where organisms learn to organise their behaviour under the influence of goals, and expected future reward is believed to drive action selection, as seen during conditioning. Neurophysiological recordings of single dopamine neurons in primates have identified a reward prediction error signal of the discrepancy between actual and expected future reward (Schultz et al., 1997; Hollerman & Schultz, 1998). In conditioning, before learning, this phasic burst of dopamine occurs at the time an unexpected reward is encountered. As trials progress and learning continues, the reward becomes more and more predictable and the phasic burst effectively moves backwards towards the time the conditioned stimulus (CS) occurs. Eventually, when full learning has taken place, the CS will elicit the same phasic response previously associated with the unexpected reward.

In particular, evidence suggests that the dopamine system may mediate the Incentive Saliency of rewards, modulating their motivational value, which is dissociable from hedonia and reward learning (Berridge & Robinson, 1998). The modern Incentive Saliency Theory distinguishes 'wanting' from 'liking,' and the dopamine system is regarded as that which calculates the 'want' rather than the 'act' parts of instrumental behaviour. Kapur's framework of psychosis builds on this hypothesis, and sees the role of dopamine as mediating the saliency of both internal and environmental representations.

Abi-Dhargham refers to the dopamine hypothesis of schizophrenia and uses neuroimaging techniques such as

SPECT and PET to monitor changes in synaptic dopamine levels. Using data from electrophysiological techniques on a smaller timescale, Winterer & Weinberger are more explicit and refer to the apparent ability of dopamine to optimise the signal to noise ratio (SNR) of local cortical microcircuits.

Carlsson et al. take a wider view and see dopamine as one of many possible dysfunctional neurotransmitters affected in the brain in schizophrenia. Pharmacological evidence suggests small differences in the fragile balance between multi-neurotransmitters at various points in local cortical microcircuits leads to many of both the positive and negative symptoms associated with the disorder. They posit that although there may be an elevated baseline release of dopamine in schizophrenia, it is possibly secondary to hypoglutamatergia.

Where is the Dysfunction?

One of the few biological disturbances that have been identified in schizophrenic patients is an impaired dopamine system, which traditionally has been of an increase in dopamine signaling in the striatum, leading to psychosis (Winterer & Weinberger, 2004). The Dopamine Hypothesis of Schizophrenia arose as a result of two major findings: (i) Exposure to dopamine receptor agonists, such as amphetamine, induces psychosis, and (ii) antipsychotic drugs provide an antipsychotic effect by blocking dopamine receptors (Abi-Dhargham, 2004). Current views still posit deficits due to an increase in dopamine; however it is the site of the excess that is controversial. Kapur refers to a general excess, while Abi-Dhargham refers to the traditional cortical/subcortical imbalance, with an excess in the subcortex and a deficit in the cortex. Winterer & Weinberger, on the other hand, suggest that it may be the cortical and not striatal microcircuits that give rise to abnormal dopamine signaling. Carlsson et al. also refer to possible cortical steering of subcortical systems, but by glutamate action. However, all agree that it is the resulting imbalance that leads to the problem and, overall, current opinion would imply that it is the imbalance in the dopamine circuits between cortical and striatal brain regions that leads to the dysfunction, while the actual point of the dysfunction remains controversial. Indeed it may be that disruption at different points in the circuits may lead to different symptoms or cognitive deficits and computational modelling may help us to answer these questions.

Carlsson et al. refer to a secondary general elevated baseline release of dopamine in schizophrenia, possibly due to a primary disturbance in cortical glutamate/GABA mediated steering of monoamine subcortical systems, (including dopamine). There is a direct glutamate pathway which acts as an accelerator and an indirect glutamate pathway that activates GABA and is an effective brake on the activity of monoamines. It is the balance between accelerator and brake that maintains stability and glutamatergic failure in the cerebral cortex may lead to negative symptoms, while glutamatergic failure in the basal ganglia would favour

positive symptoms. These result from dysregulation of the dopamine system.

Abi-Dhargham and Winterer & Weinberger also refer to such an insufficient inhibitory brake as the possible nature of the dysfunction. Abi-Dhargham refers to a hypostimulation in the cortex of D1 receptors which causes a deficit in working memory, and a hyperstimulation in the subcortex of D2 receptors which leads to psychotic symptoms, as a result of the reduced cortical brake. Winterer & Weinberger refer principally to a reduced prefrontal dopamine D1/D2 receptor activation ratio which leads to a lower cortical SNR. They posit that normally it is the D1 receptors that dominate, but in schizophrenia D2 receptors dominate, and as a result of the primary disturbance, secondary effects will occur subcortically in the striatum leading to contextually inappropriate, inflexible and bizarre behavioural routines.

All these theories seem to point to an imbalance in the dopamine system between the cortical and subcortical areas, due to an insufficient inhibitory brake system, with negative symptoms occurring as a result of disturbance to the cortex and positive symptoms as a disturbance to subcortical areas.

What is the Dysfunction?

Kapur posits that psychosis is a state of aberrant salience, where excess levels of dopamine are no longer stimulus-linked and context-driven. Delusions (paranoia, aliens interfering with one's brain), and hallucinations (hearing voices), may arise then as a result of the individual attempting to provide their own explanations for experiences which come out of the blue and are imbued with high importance. This is in keeping with an earlier theory of schizophrenia by Maher (1988) that patients make normal attributions, or reasoned normally to abnormal experiences, i.e., subcortical abnormality with normal cortical function. It is known that patients with schizophrenia suffer from a wide-spread cognitive dysfunction that affects memory, executive functioning and attention (Bilder et al., 2000; McKenna, 1997). However, there seems to be a dissociation between the psychotic experiences (delusions, hallucinations) and cognitive dysfunction. The latter occur well in advance of onset symptoms, and the trajectory of symptom recovery is not matched by cognitive recovery (Harvey, Koren, Reichenberg & Bowie, 2005). Traditional cognitive models of schizophrenia based on cognitive dysfunction in memory/attention/executive dysfunction have poor face validity when used to explain the spontaneous experiences (delusions/hallucinations) which are bizarre, or strange, since these are unrelated to past experience and stored memories (Simpson, Done, Valeé-Tourangeau, 2002).

The recent developments in understanding the role of dopamine in salience allocation do permit the formulation of cognitive neuroscience models which can integrate both Maher's theory with the known cognitive dysfunctions in schizophrenia. Computing the salience of stimuli (both external and internal, such as thoughts/ideas) is probably achieved by midbrain/ventral striatal dopamine systems rather than cortical ones (O'Doherty Dayan, Schultz, Deichmann,

Friston & Dolan, 2004). This is the ‘critic’ in models of the dopamine system in the basal ganglia (Montague, Hyman & Cohen, 2004; Sutton & Barto, 1998). In schizophrenia we posit that within the critic, the signal (winner) is distinguished from the noise (losers). This signal is then transmitted to other systems (e.g., ‘actor’ in dorsal striatum), or cortical systems, such as the dorsolateral prefrontal cortex (DLPFC) responsible for various attributional, memory, executive and attentional processes (Durstewitz, Kelc & Gunturkun, 1999). Thus stimuli, or experienced episodes, which are unimportant, can be imbued with a high degree of salience by the critic in the ventral striatum/midbrain. This provides the spontaneous experience imbued with importance. High variance in the level of background dopamine activity would also mean that these experiences occur from time to time, but not all of the time. Dopamine abnormalities in DLPFC would not only account for the neuropsychological deficits found in schizophrenia but they could also integrate the abnormal experiences into dysfunctional attributional, executive and memory systems. We can crudely equate these dual roles as being due to dopamine abnormalities in the midbrain/striatum and cortex respectively, as outlined previously in the theories of Abi-Dhargham (2004) and Winterer and Weinberger (2005). As described previously, the interaction between these different levels means that they cannot operate independently, but in consort. This permits a more tractable model of the psychology of schizophrenia, i.e. a model of both symptoms and classical cognitive abnormalities.

Antipsychotics

The action of antipsychotic drugs can help further understand what is going wrong with the dopamine system. Kapur proposes that antipsychotics dampen ‘aberrant saliences’ by blocking excess dopamine, leading to an attenuation of motivational salience of ideas and perceptions. In this way antipsychotics remove the degree to which symptoms occupy the mind, but not the core content of the symptom. They simply provide a neurochemical balance where dopamine levels return to normal, new aberrant saliences are less likely to form and existing ones are more likely to stop. It is only in the weeks to come that an individual may work through and resolve their delusions in their own time. In this way the delusions and hallucinations may be deconstructed, but this is not always the case as some patients are never able to resolve their symptoms psychologically.

Abi-Dhargham does not refer to antipsychotic action, but Winterer & Weinberger deviate from the traditional view of antipsychotic action on D2 receptors in the striatum and, using evidence from imaging studies, suggest that antipsychotics may exert actions instead through D2 receptor blockade in the cortex. Carlsson et al. refer to the adverse effects of classic antipsychotics which lead to hypodopaminergia in patients in remission from their positive symptoms that cause failure of the reward system leading to dysphoria and anhedonia; and negative effects, such as catatonia and cognitive deficits. They have developed both partial dopamine-receptor agonists, and antagonists, that act

on D2 receptors, stabilising the elevated dopamine levels without causing hypodopaminergia. However, they do not refer to the exact site of those receptors.

Both Carlsson et al. and Winterer & Weinberger focus on D2 receptor blockade as means of resolving the dopamine imbalance which leads to psychotic symptoms, but the exact site of impact remains unclear.

Interim Conclusions

Dopamine provides a RPE signal of the discrepancy between actual and expected future reward and it would appear to be an imbalance between cortical and subcortical microcircuits that leads to a dysfunction of the dopamine system. However, the actual point of the dysfunction remains controversial. Recently it has been suggested that it may be cortical microcircuits that give rise to abnormal dopamine signaling, with secondary downstream subcortical deficits, instead of the traditional view of a primary subcortical disturbance (Winterer & Weinberger, 2004).

It is generally agreed that the resulting imbalance may result from an insufficient inhibitory brake system leading to either a hypostimulation in the cortex of D1 receptors and a hyperstimulation in the subcortex of D2 receptors (Abi-Dhargham, 2004), or a reduced prefrontal dopamine D1/D2 receptor activation ratio, in which D2 receptors dominate, which leads primarily to a lower cortical SNR (Winterer & Weinberger, 2004). D2 receptor blockade would appear to be important in restoring the cortical/subcortical imbalance (Carlsson et al., 2001; Winterer & Weinberger, 2004).

Furthermore, positive psychotic symptoms arise from either a primary subcortical hyperstimulation of dopamine receptors (Abi-Dhargham, 2004), or secondary effects of either reduced cortical SNR on subcortical systems (Winterer & Weinberger, 2004), or cortical glutamate/GABA steering of subcortical systems (Carlsson et al., 2001). Negative symptoms and working memory deficits are thought to result from either hypostimulation of D1 receptors (Abi-Dhargham, 2004) or reduced prefrontal dopamine D1/D2 receptor activation ratio with D2 receptors dominating (Winterer & Weinberger, 2004).

Computational Models of Dopamine as a Reward Prediction/Temporal Difference Error Signal

Several computational models of the role of dopamine as a RPE have incorporated Temporal Difference (TD) Learning (Sutton, 1988), a form of Reinforcement Learning Theory, which provides an explicit method of modelling and quantifying the Reward Prediction error (Schultz et al, 1997; Hollerman & Schultz, 1998; Montague et al., 2004). Specifically, it provides a mathematical interpretation of how dopamine is thought to mediate reward-processing and reward-dependent learning, thus optimising behaviour in an environment. A class of TD models, known as actor-critic models (Sutton & Barto, 1998), have been adapted so that expected future reward is equivalent to incentive salience (McClure, Daw & Montague, 2003; Montague et al., 2004).

Here, the error signal generated is used in two ways: (i) The 'critic' - as a prediction error or learning signal used to create better estimates of future reward. (ii) The 'actor' - to bias action selection towards situations that predict the best reward.

It is possible that the same RPE is signaled from dopamine neurons in both the ventral tegmental area (VTA) and substantia nigra (SN). The signal is used in two ways depending on the route it takes, with the projections from VTA to ventral striatum as the 'critic' in TD models, associated with reward and motivation, and projections from SN to dorsal striatum as the 'actor', associated with motor control (O'Doherty et al., 2004; Daw, Niv & Dayan, 2005). The dopamine pathways are arranged in cortical/subcortical circuit loops involving prefrontal cortex (Alexander et al., 1985), and it is in the cortical areas that dopamine dysfunction is believed to have an effect on working memory.

It has also been suggested that TD Learning can help with the dynamic choice of action selection to obtain natural rewards required for survival. As well as assisting in the learning process, it has been suggested that the dopamine signal can be used in decision-making, when full learning has taken place, to bias the choice of actions that lead to better rewards in another actor/critic model by Schultz et al., (1997). When full learning has taken place the RPE will be zero and fluctuations above and below that point will provide important ongoing evaluations in the environment of salience which can be assessed quickly according to whether the fluctuations represent potential actions that are better or worse than expected. In this way an instant comparison can be made between well-learned possibilities; all that is required is a simple behaviour strategy, to choose those actions associated with increased dopaminergic activity and incentive salience, and avoid those of low salience where dopaminergic activity is decreased. In this way, a damaged dopamine system could explain why adults become slow to do things that they used to do so easily. Their ability to make these instant comparisons or to maintain context would become impaired, and lead to some of the cognitive deficits associated with schizophrenia, such as poor performance in the Wisconsin Card Sorting Test (WCST) or the 1-2-AX Test, where it is important to maintain context.

TD models have proved to be very successful in many behavioural tasks and are used extensively in robotics to enable learning and reacting to an environment. However, while they are often more efficient than other reinforcement learning algorithms (Suri & Schultz, 1999), complications may arise when unpredictable events occur, which break the learning chains constructed through prediction (O'Reilly & Frank, 2006), and this has led to some researchers who have previously used TD, seeking alternative combinations of algorithms as learning mechanism (Hazy, Frank & O'Reilly, In Press).

Evidence for Role of Dopamine as a Reward Prediction Error/Temporal Difference Signal

Functional imaging techniques have provided evidence that the RPE model of dopamine activity applies to human reward learning, and not just to primates, as seen in neurophysiological recordings by Schultz and colleagues mentioned above. Transient learning-related changes associated with the 'critic' have been identified in the brains of humans subjected to classical conditioning procedures, in the ventral striatum (putamen) (McClure, Berns & Montague, 2003; O'Doherty, Dayan, Friston, Critchley & Dolan, 2003). While O'Doherty et al, (2004) showed that activity in the dorsal striatum is associated with the 'actor' only, as no activity was seen in this area unless an action was required.

In addition, activation patterns consistent with predictions from a TD model of learning have also been recorded in the orbital frontal cortex (O'Doherty et al., 2003). Furthermore, Seymour et al. (2004) have used fMRI to show that neural activity in the ventral striatum and the anterior insula corresponds to the signals for sequential learning predicted by TD models, in humans in higher-order learning.

How Do Existing Computational Models Compare with the Cortical/Subcortical Debate of Theories for Schizophrenia?

The early connectionist model by Cohen & Servan-Schreiber (1992) and some biophysically detailed neural network models (Brunel & Wang, 2001; Durstewitz et al., 1999; Durstewitz, Seamans & Sejnowski, 2000) have modelled dopamine as a neuromodulator crucial for optimising the SNR thought to enhance working memory. This model is limited as it simulates only the DLPFC circuits, but not the critic in striatum and midbrain. Other models have incorporated Reinforcement Learning methods and modelled dopamine as a RPE signal, which can be effectively modelled using TD Learning (Braver et al., 1999; Suri & Schultz, 1999; Rougier et al., 2005).

As previously mentioned, it is believed that the actual point of dysfunction in subcortical/cortical microcircuits remains controversial. Cohen and colleagues have modelled working memory deficits, simulating the continuous performance test (CPT) (Cohen & Servan-Schreiber, 1992; Braver et al., 1999), with the latter model, a more powerful and complete theory of the mechanism of cognitive control, incorporating both TD learning and gating functions for dopamine, where dopamine was seen as a unitary function which enabled an organism to predict and respond appropriately to events that led to reward. In this later model schizophrenia was seen as an impaired ability to internally represent, maintain and update context relating to working memory from increased noise in the dopamine system, focusing particularly on the prefrontal cortex. The model suggested that reduced phasic activity, i.e., reduced update to active memory, led to perseveratory behaviour; while increased phasic activity, i.e., increased update, led to poor interference control, and therefore distractibility. Additionally, increased tonic (or longer-term background) activity led to delay related decay of

active memory, and therefore maintenance deficits. Both perseverations and distractibility are known disturbances to the prefrontal cortex and are typical symptoms of Schizophrenia, along with poor maintenance control. Perseveratory behaviour occurs when a patient becomes preoccupied with a task and is unable to change strategy or appropriately update goal representations, while distractibility is the inability to concentrate or focus on the task at hand. This model posits that both perseverations and distractibility are due to impairments in phasic dopaminergic activity which affect working memory. However, the model is of two very different systems in the brain doing different jobs and possibly coding for two different things; salience in the midbrain and how it possibly affects working memory in the prefrontal cortex. It is important, therefore, to investigate how these behaviours relate to each other and it is this interaction that will be explored in the current research.

The increased noise could be due to an imbalance between cortical and subcortical structures due to the insufficient inhibitory brake system on the dopamine system. However, the model has a simple architecture with no hidden layers and modules containing between one and four neurons. The simple task is hard-wired and it is not a cognitive model.

Using a more sophisticated architecture, a neural network model by Suri & Schultz (1999) specifically modelled Wolfram Schultz's work on the response of dopamine neurons in the striatum to reward-related stimuli using a 'critic', which computed and sent a TD error to an 'actor', which governed behaviour. The model did not refer explicitly to the prefrontal cortex, but showed that a reinforcement signal without RPE led to perseverations, and sustained reductions of reinforcement signal led to a loss of learned behaviour as seen in Parkinson's disease and lesioned animals.

O'Reilly and colleagues have produced a range of biophysically detailed cognitive connectionist models using O'Reilly and Munakata's Leabra algorithm, which combines error-driven and Hebbian learning with k-Winners-Take-All inhibitory competition (O'Reilly & Munakata, 2000). These models are capable of implementing the learning and gating ideas of Braver et al. (1999) mentioned above, incorporating a brake and accelerator system. A model of dynamic DA modulation in the basal ganglia by Frank (2005) separates out the roles of the D1 and D2 receptors applicable to Parkinson's disease, without using TD. The XT model (Rougier et al., 2005) uses an adaptive gating mechanism, based on an adaptive critic unit, driven by TD Learning and relates specifically to how the biological mechanisms of the prefrontal cortex support flexible cognitive control. Dorsal Lateral Prefrontal Cortex lesions were simulated by removing units and asymmetric training, resulting in perseverations in prefrontal cortex, as seen in the WCST and Stroop tasks. However, in this and all previous models, it was necessary for the dynamic gating of the basal ganglia to be hard-wired. The Prefrontal Basal Ganglia Working Memory model of learning (O'Reilly & Frank, 2006) incorporates the dynamic interactions between the prefrontal cortex and the basal

ganglia in working memory, and in doing so, abandons the use of TD in favour of an alternate associative Pavlovian mechanism. Here dopamine signals reward association and not reward prediction. Instead of using TD prediction chains over successive time-steps, which they claim break down when modelling complicated tasks such as the 1-2-AX task, the new algorithm uses the Rescorla-Wagner/Delta-rule algorithm trained by the unconditioned stimulus for the current time-step. However, this model is of learning and has not been used to model dysfunction so far.

Conclusions

The following important questions arising from recent dopamine theories of schizophrenia that remain to be addressed by current computational models:

1. Is it the cortical microcircuits that give rise to abnormal dopamine signaling with secondary downstream subcortical deficits (Winterer & Weinberger, 2004) or the traditional view of a primary subcortical disturbance?

2. For the most part connectionist models to date do not differentiate between D1 and D2 dopamine receptors, or locate the point(s) of dysfunction in the local microcircuits that give rise to a possible cortical/subcortical imbalance. They do not distinguish between the theories of Abi-Dhargham and Winterer & Weinberger, of either: (i) A hypostimulation in the cortex of D1 receptors and a hyperstimulation in the subcortex of D2 receptors (Abi-Dhargham, 2004), or (ii) A reduced prefrontal dopamine D1/D2 receptor activation ratio, in which D2 receptors dominate, leading primarily to a lower cortical SNR (Winterer & Weinberger, 2004).

3. Do positive psychotic symptoms arise from either: (i) A primary subcortical hyperstimulation of dopamine receptors (Abi-Dhargham, 2004)? Or (ii) Secondary effects of either reduced cortical SNR on subcortical systems (Winterer & Weinberger, 2004) or cortical glutamate/GABA steering of subcortical systems (Carlsson et al., 2001)?

4. Do negative symptoms and working memory deficits result from either: (i) Hypostimulation of D1 receptors (Abi-Dhargham, 2004)? Or (ii) Reduced prefrontal dopamine D1/D2 receptor activation ratio with D2 receptors dominating (Winterer & Weinberger, 2004)?

Furthermore, while enormous progress has been made regarding flexible, self-organising cognitive control, without the need for a homunculus (Rougier et al., 2005; O'Reilly & Frank, 2006), it remains to be seen whether it is prudent to abandon TD Learning, which has been shown to be an effective model of RPE (see above), or whether the problems in the break down of chaining can be overcome by some other means.

References

- Abi-Dhargham, A. (2004) Do we still believe in the dopamine hypothesis? New data bring new evidence *Neuropsychopharmacology*, 7 (supplement 1), S1-S5.
- Alexander, G. E., & DeLong, M. R. (1985) Microstimulation of the primate neostriatum. Somatotopic organization of

- striatal microexcitable zones and their relation to neuronal response properties. *J. Neurophysiol.* 53: 1417–1430.
- Berridge, K.C., & Robinson, T.E. (1998) What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Research Reviews* 28, 309-369.
- Bilder et al. (2000) Neuropsychology of first-episode schizophrenia: initial characterization and clinical correlates. *Am. J. Psychiatry*, 157(4), 549-559.
- Braver, T.S., Barch, D.M., & Cohen, J.D. (1999) Cognition and control in schizophrenia: A computational model of dopamine and prefrontal function. *Biological Psychiatry*, 46, 312-328.
- Brunel, N., & Wang, X. (2001) Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Computational Neuroscience*, 11, 63-85.
- Carlsson, A., Waters, N., Holm-Waters, S., Tedroff, J., Nilsson, M., & Carlsson, M.L. (2001) Interactions between monoamines, glutamate and GABA in Schizophrenia. *Ann. Rev. Pharm. Toxicol.* 41(237), 237-60.
- Cohen, J.D., & Servan-Schreiber, D. (1992) Context, Cortex, and Dopamine: A Connectionist Approach to Behavior and Biology in Schizophrenia. *Psychological Review* 1, 45-77.
- Daw, N.D., Niv, Y., & Dayan, P. (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioural control. *Nature Neuroscience*, 8(12), 1704-1711.
- Durstewitz, D., Kelc, M., & Gunturkun, O. (1999) A neurocomputational theory of the dopaminergic modulation of working memory functions. *Journal of Neuroscience*, 19(7), 2807-2822.
- Durstewitz, D., Seamans, J.K., & Sejnowski, T.J. (2000) Dopamine mediated stabilization of delay-period activity in a network model of prefrontal cortex. *J. Neurophysiology*, 83, 1733-1750.
- Frank, M.J., (2005) Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism. *J. of Cognitive Neuroscience*, 17, 51-72.
- Harvey, P.D., Koren, D., Reichenberg, A., & Bowie, C.R. (2005) Negative Symptoms and Cognitive Deficits: What is the nature of their Relationship?, *Schizophrenia Bulletin*, (Advanced access 12/10/05).
- Hazy, T.E., Frank, M.J., & O'Reilly, R.C (In Press) Banishing the homunculus: Making working memory work. *Neuroscience*.
- Hollerman, J.R., & Schultz, W. (1998) Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, 1, 304-309.
- Kapur, S., (2003) Psychosis as a state of aberrant salience: A framework linking biology, phenomenology, and pharmacology in schizophrenia. *American Journal of Psychiatry*, 160(1), 13-23.
- McClure, S.M., Berns, G.S., & Montague, P.R. (2003) Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, 38, 339-346.
- McClure, S.M., Daw, N.D., & Montague, P.R., (2003) A computational substrate for incentive salience. *Trends in Neuroscience* 26(8), 423-428.
- McKenna, P.J. (1997) *Schizophrenia and related syndromes*. Hove: Psychology Press.
- Maher, B.A. (1988) Anomalous experience and delusional thinking: The logic of explanations. In T.F.Oltmanns and B.A. Maher (Eds.), *Delusional beliefs* (pp. 15-33), New York: Wiley.
- Montague, P.R., Hyman, S.E., & Cohen, J.D. (2004) Computational roles for dopamine in behavioral control. *Nature* 431, 760-767.
- O'Doherty, J.P., Dayan, P., Friston, K., Critchley, H., & Dolan, R.J. (2003) Temporal difference models and reward-related learning in the human brain. *Neuron*, 38, 329-337.
- O'Doherty, J.P., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R.J (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, 304, 452-454.
- O'Reilly, R.C., & Frank, M.J. (2006) Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation* vol 18(2), 283-328.
- O'Reilly, R.C., & Munakata, Y. (2000) *Computational Explorations in Cognitive Neuroscience: Understanding the mind by simulating the brain*. MIT Press.
- Rougier, N.P., Noelle, D.C., Braver, T.S., Cohen, J.D., & O'Reilly, R.C. (2005) Prefrontal cortex and flexible cognitive control: Rules without symbols. *PNAS*, 102(20), 7738-7343.
- Schultz, W. (1998) Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80, 1-27.
- Schultz, W., Dayan, P., Montague, P.R. (1997) A Neural substrate of prediction and reward. *Science*, 275: 5306, 1593-1599.
- Seymour et al. (2004) Temporal difference models describe higher-order learning in humans. *Nature*. 429, 664-667.
- Simpson J, Done D.J., Valeé-Tourangeau (1998) An Unreasoned Approach: A Critique of Research on Reasoning and Delusions. *Cog. Neuropsychiatry*, 3, 1-20.
- Suri, R.E., & Schultz, W. (1999) A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neurosci.*, 91(3), 871-890.
- Sutton, R.S. (1988) Learning to Predict by the Methods of Temporal Differences. *Machine Learning* 3, 9-44.
- Sutton, R.S., & Barto, A.G. (1998) *Reinforcement learning*, MIT Press.
- Winterer, G., & Weinberger, D.R. (2004) Genes, dopamine and cortical signal-to-noise ratio in schizophrenia. *Trends in Neurosciences*, 27(11), 683-690.