

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/123415>

Please be advised that this information was generated on 2019-12-11 and may be subject to change.

Dissociable Effects of Dopamine and Serotonin on Reversal Learning

Hanneke E.M. den Ouden,^{1,3,*} Nathaniel D. Daw,^{2,3} Guillén Fernández,^{1,4} Joris A. Elshout,^{1,4} Mark Rijpkema,¹ Martine Hoogman,⁵ Barbara Franke,^{1,5,6} and Roshan Cools^{1,6}

¹Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen 6500, the Netherlands

²Center for Neural Science, New York University, New York, NY 10003, USA

³Department of Psychology, New York University, New York, NY 10003, USA

⁴Department of Cognitive Neuroscience, Radboud University Nijmegen Medical Centre, Nijmegen 6500, the Netherlands

⁵Department of Human Genetics, Radboud University Nijmegen Medical Centre, Nijmegen 6500, the Netherlands

⁶Department of Psychiatry, Radboud University Nijmegen Medical Centre, Nijmegen 6500, the Netherlands

*Correspondence: h.denouden@gmail.com

<http://dx.doi.org/10.1016/j.neuron.2013.08.030>

SUMMARY

Serotonin and dopamine are speculated to subserve motivationally opponent functions, but this hypothesis has not been directly tested. We studied the role of these neurotransmitters in probabilistic reversal learning in nearly 700 individuals as a function of two polymorphisms in the genes encoding the serotonin and dopamine transporters (*SERT*: 5HTTLPR plus rs25531; *DAT1* 3'UTR VNTR). A double dissociation was observed. The *SERT* polymorphism altered behavioral adaptation after losses, with increased lose-shift associated with L' homozygosity, while leaving unaffected perseveration after reversal. In contrast, the *DAT1* genotype affected the influence of prior choices on perseveration, while leaving lose-shifting unaltered. A model of reinforcement learning captured the dose-dependent effect of *DAT1* genotype, such that an increasing number of 9R-alleles resulted in a stronger reliance on previous experience and therefore reluctance to update learned associations. These data provide direct evidence for doubly dissociable effects of serotonin and dopamine systems.

INTRODUCTION

Dopamine and serotonin have both long been implicated in behavioral control and decision-making. One central idea is that these neurotransmitters are involved in learning from reinforcement. This theory is most strongly supported by experimental findings on dopamine, where notable progress has been made in the last two decades. Groundbreaking electrophysiological studies showed that dopaminergic neurons in the midbrain increase firing to outcomes that exceed expectations (Fiorillo et al., 2003; Schultz et al., 1997). Advances in theoretical modeling then envisioned phasic dopamine responses as a reinforcement signal, “stamping in” successful operant responses

(Frank et al., 2004; Houk et al., 1995; Montague et al., 1996; Suri and Schultz, 1999). Pharmacological and fMRI studies in humans support this idea, showing that dopaminergic drugs enhance relative learning from reward compared to punishments in both healthy individuals (Cools et al., 2009) and patients with Parkinson's disease (Cools et al., 2006; Frank et al., 2004).

Although there is no similarly well-developed theoretical or formal framework for guiding and interpreting empirical research on serotonin, serotonin has been most closely associated with learning from negative events. For example, after administration of the serotonin reuptake inhibitor citalopram, healthy subjects shift away more frequently from a stimulus that resulted in a loss (Chamberlain et al., 2006), and lowering levels of serotonin using dietary tryptophan depletion selectively improves the prediction of punishments (Cools et al., 2008b). More specifically, serotonin has been associated with the inhibition of punished behaviors (Crockett et al., 2009; Dayan and Huys, 2008; Deakin and Graeff, 1991; Soubrie, 1986). Taken together, these results support the notion that dopamine and serotonin are involved in learning from reward and punishments, respectively (although see e.g., Maia and Frank, 2011; Palminteri et al., 2012; Robinson et al., 2010). It was recently suggested that their actions are characterized by mutual opponency (Boureau and Dayan, 2011; Cools et al., 2011; Daw et al., 2002).

However, both neuromodulators have also been implicated in another key set of behaviors, namely the ability to flexibly change behavior. In order to successfully interact with our environment, it is important to be able to ignore rare events in a stable environment, yet to flexibly update our beliefs when our environment changes. Such an optimal balance of cognitive stability and flexibility depends on successful integration the consequences of our actions over a longer timescale. Perseverative behavior is the tendency to stick to a particular choice independent of, or even in spite of, contrary evidence and reflects the failure to flexibly adapt. Dopamine manipulations in both rodents and humans selectively altered behavior and neural processes associated with the ability to reverse previously rewarded choices (Boulougouris et al., 2009; Clatworthy et al., 2009; Cools et al., 2009; Dodds et al., 2008; Rutledge et al., 2009). With respect to serotonin, antagonists of the 2A and 2C receptors affected the number of errors during reversal before reaching a preset learning

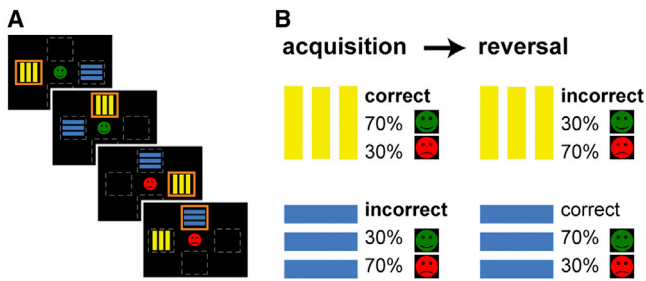


Figure 1. Probabilistic Reversal Learning Paradigm

(A) On each trial, two stimuli were presented in two out of four randomly selected locations. The subject was instructed to select the usually rewarded stimulus, and feedback was presented in the form of positive or negative emoticons. Win-stay trials were trials on which the subject picked the same stimulus as they did on the previous, rewarded trial (e.g., trial 2 and 3 in A). Lose-shift trials are trials on which the subject shifted response after a punishment (e.g., trial 4).

(B) During acquisition, the correct stimulus (here yellow) resulted in a 70:30 ratio of reward/punishment. Choosing the incorrect stimulus (blue) led to the reverse (30:70) ratio. After 40 trials, contingencies reversed and the subject had to learn to now select the blue stimulus. Any trial in which the “incorrect” stimulus was chosen was defined as an error trial.

criterion (Boulougouris et al., 2008; Boulougouris and Robbins, 2010), and serotonin depletion in the orbitofrontal cortex in nonhuman primates increased the number of perseverative errors on a deterministic reversal learning task (Clarke et al., 2007). These two functions of learning from reinforcement versus behavioral flexibility can perhaps be reconciled if we view perseveration as another manifestation of reinforcement-like effects that are accumulated during the prereversal phase. In other words, they might provide a different window on the same underlying functionality.

In the present study, we take a behavioral genetics approach to study the role of serotonin and dopamine in human decision making. The advantage of this approach over pharmacological studies is that it allowed us to test effects of both transmitter systems within single individuals, in a single session and in a large cohort. In this study, we pit against each other polymorphisms that affect dopamine and serotonin function to assess their dissociable and opponent roles in decision making. We investigate effects of polymorphisms in the regulatory regions of the serotonin and dopamine transporter genes: the *SLC6A4/SERT/5HTT*-length polymorphism (5HTTLPR) combination with a single nucleotide polymorphism within the repeat (rs25531) and a variable repeat in the 3' regulatory region of *SLC6A3/DAT1*. Although the exact functional consequences of these polymorphisms on serotonin and dopamine transmission are as yet unclear, evidence from multiple sources confirms that these polymorphisms can be used to investigate effects of the dopamine and serotonin systems. In vitro, the *DAT1* and *SERT* polymorphisms cause natural variation in the expression levels of these transporters (Hu et al., 2006; Mill et al., 2002). In addition, PET/SPECT studies in humans have shown reduced SERT binding in S'-carriers (Willeit and Praschak-Rieder, 2010) and higher striatal DAT availability in carriers of the 9-repeat (9R) allele of *DAT1* (Spencer et al., 2013; van de Giessen et al., 2009; van Dyck et al., 2005, although see Costa et al.,

2011). Furthermore, the effects of these polymorphisms on behavior and brain function as well as their association with psychiatric disorders tend to follow the functional dimensions associated with serotonin (Caspi et al., 2010; Hariri and Holmes, 2006; Lesch et al., 1996; Roiser et al., 2009) and dopamine (Aarts et al., 2010; Forbes et al., 2009; Franke et al., 2010; Gizer et al., 2009).

To independently assess the effects of serotonin and dopamine on both immediate effects of reinforcement on subsequent choices and on longer-term behavioral flexibility, we use a probabilistic reversal learning paradigm. First, to examine direct outcome reactivity, we assess the tendency to locally shift responding immediately after negative feedback and to stick to a response after positive feedback. We hypothesize that the *SERT* polymorphism will alter lose-shifting, whereas *DAT1* variation will affect win-staying. Such behavior would be a direct manifestation of reinforcement properties hypothetically associated with either neurotransmitter, as embodied in Thorndike's law of effect (Thorndike, 1911) or in computational models such as temporal difference learning.

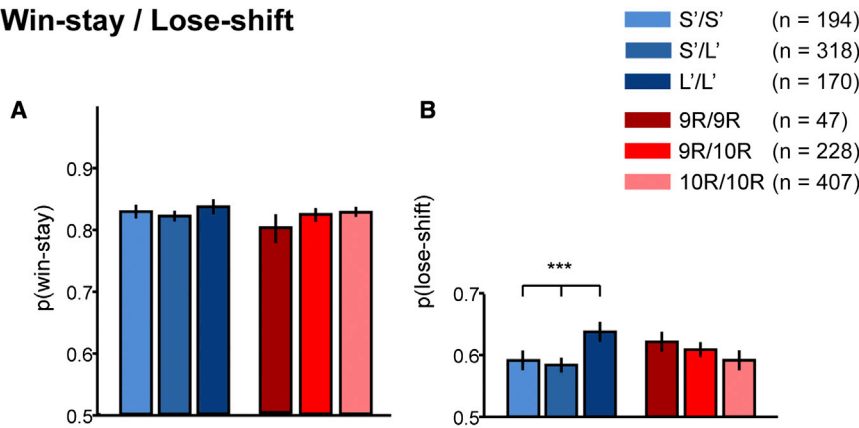
Second, we analyze the effects of the *SERT* and *DAT1* polymorphisms on choices after reversal to assess perseveration. As mentioned above, perseveration might be an additional consequence of reinforcement, separate from any more local effects on win-stay/lose-shift behavior. In a reversal task, perseveration on a previously favored alternative following reversal might reflect the repeated reinforcement of that response accumulated during the prereversal phase. If a strongly stamped-in response tendency takes repeated trials before it is unlearned, then a reinforcement mechanism such as that associated with dopamine would give rise to perseveration at time of reversal. Another possibility is that perseveration occurs due to a failure to learn from the negative feedback that now follows a previously rewarded stimulus. We compare such potential perseveration mechanisms by fitting computational learning model to our data and subsequently test whether their estimated parameters are affected by genotype.

RESULTS

Subjects ($n = 810$) completed a probabilistic reversal learning task (see Table S1, available online, for demographic information). On each trial, they selected one of two stimuli, which led probabilistically to either reward or punishment (Lawrence et al., 1999) (Figure 1). During the first 40 trials, stimulus A was usually rewarded (70%), but sometimes punished (30%), and vice versa for stimulus B. For the second 40 trials, these contingencies were reversed. Subjects were instructed to select the usually rewarded stimulus (for details see Experimental Procedures).

All subjects were genotyped for *SERT* and *DAT1* polymorphisms. Full behavioral, genetic, and demographic data were available for 685 participants, from which three subjects were excluded for failure to perform the task (for details on genotyping and exclusions see Supplemental Experimental Procedures). There was no significant difference between genotypes in gender distribution (both polymorphisms: $\chi^2(2) < 4$, $p > 0.1$).

Win-stay / Lose-shift



Perseveration

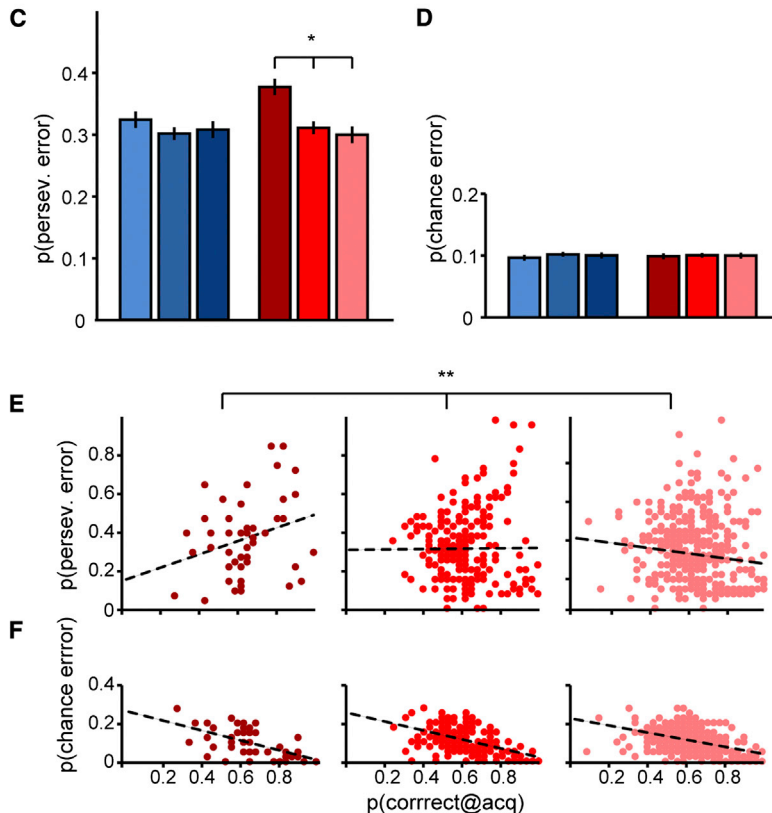


Figure 2. Win-Stay/Lose-Shift and Perseveration Results

(A–D) Win-stay/lose-shift: L'-homozygotes of the *SERT* polymorphism showed significantly more lose-shifting than S'-carriers. (A) There was no effect of *DAT1* on lose-shifting. (B) There was no effect of *SERT* or *DAT1* on win-stay. (C and D) Perseveration: a higher 9R:10R allele ratio was associated with more perseveration, whereas there was no effect of *DAT1* on chance error rate. There was no effect of *SERT* on perseverative or chance error rates. Mean \pm SEM. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.005$.

(E) There was an interaction of number of correct choices during acquisition and *DAT1* polymorphism, the relationship between choice history and perseveration reversed as a function of genotype.

(F) There was a negative effect of choice history on chance error rates, but no interaction with *DAT1*. See also Figure S1 and Tables S2 and S3.

Supplemental Experimental Procedures, Figure S1, and Table S2). Both *SERT* and *DAT1* selectively affected these three measures (*SERT*: $F(3.7, 1189) = 3.38$, $p = 0.011$, $\eta^2 = 0.010$; *DAT1*: $F(3.7, 1189) = 3.07$, $p = 0.019$, $\eta^2 = 0.09$). Below, we explore the nature of these main effects of measures of interest.

Win-Stay/Lose-Shift

Consistent with our hypothesis, *SERT* affected the likelihood of shifting responses after punishment ($F(20,661) = 5.80$, $p = 0.003$, $\eta^2 = 0.017$; Figure 2A). Pairwise post hoc comparisons revealed that L' homozygotes exhibited increased lose-shift rate relative to the S' carriers, whereas there was no difference between the S' homozygotes and the heterozygotes (L'/L' > S'/S', $p = 0.001$; L'/L' > S'/L', $p = 0.033$; S'/S' versus S'/L', $p = 0.15$). Indeed, grouping S'-carriers versus L'-homozygotes does not alter significance ($F(15,666) = 9.28$, $p = 0.002$, $\eta^2 = 0.014$). Conversely, there was no effect of *SERT* on win-stay rates (Figure 2B).

In contrast to our hypothesis, *DAT1* did not affect win-stay (or lose-shift) rates (Figures 2A and 2B). There were also no gene-gene interactions between the two polymorphisms for either win-stay or lose-shift (all $F(20,661) < 1.5$, $p > 0.3$, $\eta^2 < 0.001$). There was no effect of gender, age, or education on win-stay or lose-shift (all tests: $F(20,661) < 3$, $p > 0.1$).

As mentioned in the introduction, probabilistic discrimination and reversal tasks require subjects to ignore rare events in a stable environment, yet adjust their responses when the environment has changed. Therefore, we next assessed whether the

Probabilistic Reversal Learning

Our primary analysis focused on three main measures of interest: win-staying, lose-shifting (both as a function of the previous trial), and perseveration. Perseverative errors were defined as any sequence of two or more errors during the reversal phase. These three measures were included as within-subject measures in a repeated-measures ANOVA, together with the between-subject factors gender and learning criterion attainment, and covariates age and level of education (for control analyses of basic learning measures and covariates, see

SERT genotype affected response adaptation after any negative feedback, or whether this was specific to either the feedback validity or task epoch (acquisition or reversal). There was no interaction of *SERT* genotype with feedback validity ($F(2,668) = 0.5$, $p = 0.6$, $\eta^2 = 0.001$), and *SERT* genotype significantly affected lose-shift whether feedback was invalid ($F(2,668) = 4.8$, $p = 0.009$, $\eta^2 = 0.014$) or valid ($F(2,668) = 5.3$, $p = 0.005$, $\eta^2 = 0.016$). This is not surprising, given that subjects are not aware of feedback validity. There was also no interaction of *SERT* genotype and task phase ($F(2,668) = 1.9$, $p = 0.15$, $\eta^2 = 0.006$), and the effect of *SERT* genotype on lose-shift was significant during both the acquisition phase ($F(2,668) = 6.3$, $p = 0.002$, $\eta^2 = 0.018$) and the reversal phase ($F(2,668) = 3.1$, $p = 0.047$, $\eta^2 = 0.009$).

Perseveration

A hierarchical regression analysis showed that *DAT1* genotype significantly predicted the proportion of perseverative errors during the reversal phase, such that a higher ratio of 9R:10R alleles led to an increased number of perseverative errors ($\beta = 0.084$, $t(671) = 2.22$, $p = 0.029$) (Figure 2C). This effect was specific to perseveration, as evidenced by the finding that there was no effect of *DAT1* on chance errors ($t(671) = 0.07$, $p = 0.95$) (Figure 2D), which were defined as single errors that occurred between two correct responses.

Furthermore, there was an effect of *DAT1* genotype on the interaction between perseveration and the choice history (rate of correct responses during acquisition; $\beta = 0.10$, $t(671) = 2.72$, $p = 0.007$) (Figure 2E), in the absence of a main effect of choice history on perseverative error rate ($t(671) = 0.44$, $p = 0.66$). Again, there was no such interaction for chance errors ($t(671) = 1.5$, $p = 0.14$).

The *DAT1* effects of choice history on perseveration were characterized by a dose-dependent reversal of their relationship: in 9R homozygotes perseveration increased with increasing number of correct choices during acquisition ($\beta = -0.34$, $t(40) = 2.6$, $p = 0.013$), whereas in heterozygotes there was no association ($\beta = 0.061$, $t(221) = 0.89$, $p = 0.38$), and in 10R homozygotes perseveration marginally decreased ($\beta = -0.092$, $t(400) = -1.8$, $p = 0.069$). We verified this effect against sensitivity to outliers using a robust regression, which confirmed the dose-response effects (9R9R, $\beta = 0.062$, $t(40) = 2.31$, $p = 0.026$; 9R10R, $\beta = -0.008$, $t(221) = -0.61$, $p = 0.54$; 10R10R: $\beta = -0.024$, $t(400) = -2.7$, $p = 0.007$).

The *SERT* genotype did not affect any type of reversal errors ($p > 0.5$) (Figures 2C and 2D). In addition, sex, age, or education covariates did not explain a significant proportion of variance in any of the reversal error scores ($R^2 < 0.01$, $F(3,678) < 1.8$; $p > 0.1$).

In summary, the present data set reveals a double dissociation between effects of the *SERT* and *DAT1* genotypes on reversal learning, with *SERT* altering global lose-shifting and *DAT1* altering postreversal perseveration. In a final ANOVA, we ascertained that the relative difference in lose-shift and perseveration *Z* scores was predicted by the difference in *SERT* and *DAT1* genotype ($R^2 = 0.16$, $F(5,676) < 25.5$; $p = 0.009$). This significant interaction confirms the double dissociation between the two effects, with *SERT* affecting lose-shifting but not perseveration, and *DAT1* affecting perseveration but not lose-shifting.

Computational Model

We next used computational models to investigate the mechanisms that might underlie the *DAT1* genotype results. Although *DAT1* shows robust effects in our data set, the measure of perseveration to which it is related is relatively opaque, in contrast to the more direct measure of trial-by-trial switching with which *SERT* was associated.

This opaqueness results from the fact that (perseveration) error scores require some form of “topdown” definition or knowledge by the experimenter, e.g., when the reversal, unbeknownst to the subject, has occurred. This has hampered comparison of previous studies of reversal learning studies, which have reported a veritable zoo of reversal error measures, such as errors to criterion, total reversal errors, maintenance errors, perseverative errors, learning errors, and chance errors. Models of reinforcement learning can provide a more principled approach to assessing behavior, because they are independent of such external definitions that the subject is unaware of (learning criterion, point of reversal). Instead, like for win-stay/lose-shift measures, they take into account only past choices and observed outcomes.

We aimed to understand the process or mechanism underlying the effect of *DAT1* on perseveration using a reinforcement learning model to examine how perseveration can arise from a learning process integrating reward over a longer timescale. For simplicity, we do not consider the more transparent *SERT* effects on lose-shift behavior here, although we have verified in simulations not reported here that our model captures them when it is augmented with an additional parameter that directly controls switching after losses, without affecting long-term value integration.

In the context of reinforcement learning models, two features of the *DAT1* effects are puzzling. First, the effect is selective to the reversal phase, and second, the relationship between performance in the acquisition and reversal phases reverses sign depending on genotype. A standard account such as a temporal-difference learning model predicts neither of these features because learning during both phases is driven by a common mechanism. Unaugmented, such a model predicts that errors on either phase should track one another. In particular, the learning rate parameter affects the acquisition and reversal equally, by speeding up or slowing down acquisition and updating of associations. The inverse temperature parameter also affects errors in both phases equally, where a decrease will lead to lead to more random (i.e., less value-driven) choices globally.

Accordingly, we considered a model that generalizes temporal-difference learning to include an “experience” weight parameter (ρ), which decouples acquisition and reversal by allowing the balance between past experience and new information to increasingly tip in favor of past experience. This feature is derived from the experience-weighted attraction (EWA) model (Camerer and Ho, 1999), although we do not include additional features from that model that relate to its use in modeling multiplayer games. The action of the experience weight parameter captures the intuition that reinforcement accumulated over the course of the acquisition phase could make it relatively more difficult to adjust when the contingencies are reversed, leading to perseveration. The experience weight

Table 1. Model Parameters Used to Fit the EWA and RP Models

Model	Parameter	Prior	Constraint	Median	Range (25%–75%)	x_{p_i}	$p(m_i data, m)$
EWA							
	ϕ	β (1.2,1.2)	$0 < \phi < 1$	0.33	0.10–0.62	1	0.694
	ρ	β (1.2,1.2)	$0 < \rho < 1$	0.62	0.28–0.88		
	β	Gaussian (0,10)	$-\infty \leq \beta \leq \infty$	4.69	2.62–7.35		
RP							
	α^{rew}	β (1.2,1.2)	$0 < \alpha^{rew} < 1$	0.73	0.36–0.93	0	0.305
	α^{pun}	β (1.2,1.2)	$0 < \alpha^{pun} < 1$	0.63	0.25–0.84		
	β	Gaussian (0,10)	$-\infty \leq \beta \leq \infty$	4.22	2.13–7.65		

For each parameter we used weakly informative priors. Median and range of the fitted parameters across all subjects. x_{p_i} , exceedance probability for model i .

parameter interpolates between a standard temporal-difference learning model ($\rho = 0$), where predictions are always driven by the most recent experiences, and a model ($\rho = 1$) that weights all trials in the experiment equally, causing all the experience accumulated during the acquisition phase to produce sluggish reversal.

For comparison, we tested a more standard reinforcement learning model to determine whether the experience weight parameter is superior in capturing behavioral strategies and genotypic effects. This model is also based on the classic Rescorla-Wagner model of conditioning, but in this case, expanded with separate learning rates for reward (α^{rew}) and punishment (α^{pun}) trials (“RP model”) (Frank et al., 2007). If *DAT1* were selectively related to (α^{pun}), then this might provide a different explanation for the gene’s selective relationship to perseveration following reversal, if errors during acquisition relate more to positive feedback and during reversal to negative feedback. In particular, if the string of punishments observed immediately after reversal has little effect, then it will take longer to update the value of the chosen stimulus.

Model Comparison and Parameter Inference

After fitting both models on a trial-by-trial basis to each individual, Bayesian model comparison showed that the EWA model was superior to the RP model (Table 1, exceedance probability = 1.00).

Next, we used the estimated model parameters from the winning EWA model to simulate choices. This cycle of fitting and resimulation allowed us to analyze these simulated choices in the same way we analyzed the original data to assess whether the fitted model is able to capture the observed differences as a function of *DAT1* genotype, and if so, how. First we visualized the overall learning curves of the simulated subjects. Figure 3A shows the trial-by-trial estimated probability of choosing the stimulus that was correct (i.e., 70% rewarded) during acquisition and incorrect during reversal. This figure confirms that the model captures the differential effects of *DAT1* on perseveration in the absence of any differences during acquisition. With an increasing number of 9R alleles, the simulated subjects are more likely to perseverate, i.e., more likely to choose the originally correct stimulus during reversal.

We subsequently analyzed the choices simulated by the model in the same manner as the original data. Using the fitted parameters, the model replicated all the *DAT1*-related behaviors

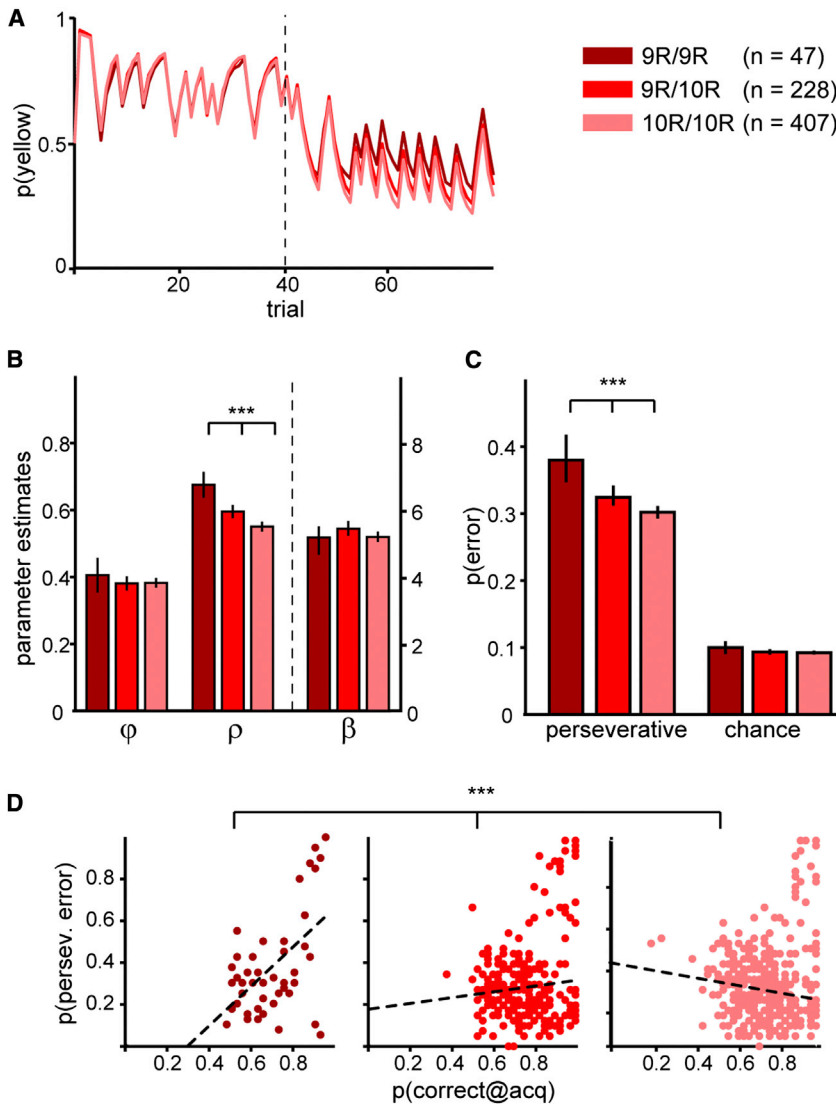
shown by our participants. There was a significant main effect of *DAT1* on the perseverative error rate (Figure 3C) ($\beta = -0.02$, $t(671) = -2.7$, $p = 0.007$), in the absence of such an effect on the chance error rate ($t(671) = -0.48$, $p = 0.6$) or on win-stay or lose-shift rates (both: $F(17,664) < 1$, $p > 0.5$, $\eta^2 < 0.002$). In addition, the model also captured the dose-dependent reversal of the effect of the choice history on perseveration (Figure 3D) (*DAT1* × choice history: $t(671) = 4.9$, $p < 0.001$; 9R9R, $\beta = 0.144$, $t(40) = 4.4$, $p < 0.001$; 9R10R, $\beta = 0.009$, $t(221) = 0.74$, $p = 0.46$; 10R10R: $\beta = -0.024$, $t(400) = -3.22$, $p = 0.001$).

To understand what features of the model were producing the behavioral effects, we examined how the best fitting parameters varied with genotype. Jonckheere’s test revealed that the experience weight ρ significantly increased with the number of 9R alleles ($J = 53,943$, $Z = -2.88$, $p = 0.004$) (Figure 3B), in absence of any gene-dose-dependent effects on the other parameters (β : $J = 60,179$, $Z = -0.44$, $p = 0.7$; ϕ : $J = 61,542$, $Z = 0.09$, $p = 0.9$).

Finally, we conducted two control analyses on simulated data and model parameters. First, we found no significant effects of *SERT* genotype on the three parameters of the EWA model (Mann-Whitney U on L-homozygotes versus S’-carriers; β : $U = 42,147$, $Z = -0.6$, $p = 0.5$; ϕ : $U = 40,911$, $Z = -1.2$, $p = 0.24$; ρ : $U = 42,214$, $Z = -0.6$, $p = 0.6$; see also Figure S2). Second, we established that there were no significant effects of *DAT1* genotype in the RP model on reward or punishment learning rates, or a difference between these two. There were no effects of *DAT1* on any of the parameters. (α^{pun} : $J = 61,372$, $Z = 0.02$, $p = 0.9$; α^{rew} : $J = 63,672$, $Z = 0.91$, $p = 0.4$; $\alpha^{rew} - \alpha^{pun}$: $J = 63,038$, $Z = 0.67$, $p = 0.5$; β : $J = 60,417$, $Z = -0.35$, $p = 0.7$).

DISCUSSION

The present study revealed a double dissociation between serotonin and dopamine influences on reinforcement learning by comparing the effects of genetic polymorphisms in *SERT* and *DAT1*. We show that the *SERT* polymorphism selectively affects immediate lose-shift behavior, whereas variation in the *DAT1* polymorphism alters perseveration in the reversal phase. It is important to keep in mind the interpretational difficulty in terms of the direction of these effects, given that the exact functional consequences of the *SERT* and *DAT1* polymorphisms as yet unclear. Nonetheless, these findings speak against a directly



opponent role of serotonin and dopamine and rather point to differential processes of action/outcome integration that take effect on a different timescale.

SERT, Serotonin, and Lose-Shift

Allelic variation in *SERT* predicted the likelihood of behavioral adaptation after punishment but not reward. This effect was not specific to either the validity of the feedback or the phase of the task, indicating that it was a global effect on behavioral adaptation after negative feedback. The increased tendency to shift responses after punishment in *L'*-homozygotes without influencing behavior following reward is in line with opponency models that suggest a specific role for serotonin in behavioral adaptation in the face of punishment (Cools et al., 2011; Daw et al., 2002). *L'*-homozygotes have been shown to exhibit increased *SERT* binding (Willeit and Praschak-Rieder, 2010), which might lead to decreased levels of extrasynaptic serotonin. If this is the case, our results echo findings of enhanced lose-shift behavior after decreased brain serotonin levels, either by exper-

Figure 3. EWA Model Simulation

The choices simulated by the model using the estimated parameters replicate the observed *DAT1* effects in perseveration.

(A) Trial-by-trial estimated probability of choosing the initially correct stimulus, averaged by *DAT1* genotype. There is no difference during acquisition, but at reversal, an increasing 9R:10R ratio leads to sustained choice of the initially correct stimulus. i.e., perseveration.

(B) This effect is mediated by a selective change in the experience weight parameter ρ , in the absence of a change in any of the other two parameters. Importantly, ρ determines how quickly the weight of past experiences increases Mean \pm SEM. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.005$.

(C and D) The model simulations replicate the dose-dependent effect of the *DAT1* polymorphism on the perseverative error rate in absence of an effect on chance error rate (C), as well as the interaction with the choice history (D) (compare Figures 2C–2E). (C) Mean \pm SEM. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.005$. See also Figure S2 for *SERT*.

imental manipulation (Bari et al., 2010; Chamberlain et al., 2006) or as a consequence of hypothesized reductions in depression (Murphy et al., 2003). They also agree with the enhanced punishment prediction observed after tryptophan depletion, which lowers central serotonin levels (Cools et al., 2008b). The present results disambiguate contradictory effects in previous reversal learning studies with smaller sample sizes (Izquierdo et al., 2007; Jedema et al., 2010; Vallender et al., 2009), confirming a clear role for *SERT* in immediate behavioral adaptation after losses. Note that the general nature of this effect explains why there are no

global differences in task performance between the different *SERT* genotypes: although *L'*-homozygotes were more likely to choose the incorrect stimulus after a probabilistic punishment, they were also more likely to switch to the correct stimulus after a punished incorrect choice.

There was no evidence for an influence of *SERT* on the reversal aspect of the task, in contrast to previous neurochemical studies with nonhuman primates (Clarke et al., 2007; Walker et al., 2009). This discrepancy may reflect differential degrees of serotonin depletion in the different studies: serotonin depletion with the neurotoxin 5,7-DHT in marmosets produces very severe depletion, in contrast to the presumably subtle differences in baseline serotonin levels through genetic polymorphisms. Such different manipulations may well have qualitatively different effects on for example tonic versus phasic firing (Cools et al., 2008a).

DAT1, Dopamine, and Perseveration

DAT1 allelic variation specifically affected performance during the reversal phase, in the absence of any differences during

acquisition. During reversal, a higher 9R:10R allele ratio led to both an overall increase in perseverative errors, as well as a change in the influence of the choice history on perseveration. Importantly, there was no overall difference between *DAT1* genotypes in terms of acquisition scores. This dissociation between acquisition and reversal is difficult to capture in standard computational models of error-driven learning, which essentially describe a local (although incremental) win-stay/lose-shift preference adjustment mechanism by which both initial acquisition and its reversal proceed equivalently.

We were able to explain these effects on perseveration and its interaction with choice history in such a model by including an additional feature derived from the experience-weighted attraction model (Camerer and Ho, 1999). In this model, the relative weight of past experience with respect to incoming information increased every time a particular action was selected, which produced an increased reliance on current beliefs over new information. The rate of increase was determined by the experience weight decay parameter ρ . From the fitted model parameters, it appeared that *DAT1* allelic variation selectively affected the size of the experience weight decay, such that the parameter increased with an increasing number of 9R alleles. This increase resulted in a larger weight of past experience at the time of reversal for stimuli that had often been chosen, which made subjects more reluctant to update the strongly held belief about the previously rewarded stimulus, causing perseveration. Computationally, this effect can be understood as a learning rate that declines more rapidly with experience, as in uncertainty-based learning models such as the Kalman filter (Dayan et al., 2000). However, perhaps more closely related to notions of dopamine as a reinforcement signal, it can conversely be understood as an increasing tendency for previous learning to accumulate rather than decay, progressively overshadowing new learning. This may embody an aspect of the colloquial notion of reinforcement “stamping in” choices that standard temporal difference models fail to capture. Interestingly, there was no similar effect of *DAT1* genotype on overall win-stay behavior. This observation suggests that the *DAT1* variants do not affect local choice adjustment per se. Perseveration and win-stay rates both seem to represent indices of the strength of reinforcement, in the first case measured by difficulty reversing the learned knowledge, and in the latter by the immediate effect on subsequent trials. Although these two effects are coupled by a single learning mechanism in standard models, they are dissociated in our data. A crucial difference is that the win-stay rate is a local measure of the effect of reward only one trial back in time, whereas perseveration is by definition a measure of their longer-term cumulative effects. This dissociation may also relate to (dorsolateral) striatal dopamine’s hypothesized role in habitual behavior (Balleine and O’Doherty, 2010; Daw et al., 2005; Everitt and Robbins, 2005), and to the idea that in humans, local choice adjustments (e.g., win-stay) in choice tasks of this sort relate more to an explicit, working memory-based mechanism that may mask underlying incremental reinforcement learning (Collins and Frank, 2012).

Another possibility about how local adjustment and reversal might relate is that a deficit in learning from punishments (relative to reward) might exhibit itself as an apparently selective difficulty at reversal time, when a cluster of negative feedback occurs.

However, although this mechanism might predict dissociation between errors in reversal versus overall errors in initial acquisition, it does not seem to provide a good explanation of the observed pattern of *DAT1* effects. This is because such a mechanism would couple the reversal deficit with global lose-shifting, and these are doubly dissociated by our *DAT1* and *SERT* effects. Accordingly, the EWA model also provided a better overall fit to choices than an alternative model involving differential learning from reward and punishments.

An important interpretational caveat with the present study is that the task has only two, mutually exclusive response options, which makes it difficult to distinguish to what extent choice of either option relates to its own perceived strength versus the weakness of the other. For instance, it may not be definitively possible to disentangle truly perseverative responding (in the sense of a sustained affirmative tendency to seek the previously reinforced option) from impairment in acquiring or sustaining a response to the newly highly reinforced option. Nevertheless, the best-fitting model here suggests that *DAT1*-related perseveration occurs due to large, sustained value on the previously favored option. Future studies should test this model using a task with a third option.

Notwithstanding these finer distinctions, our finding relating *DAT1* to reinforcement is in line with the conditioning literature suggesting that dopamine potentiates responding to cues previously associated with reward. Specifically, studies in rodents (Goto and Grace, 2005; Parkinson et al., 1999) have shown that enhanced levels of dopamine potentiate responding to previously rewarded stimuli. Furthermore, dopaminergic medication in patients with Parkinson’s disease has been shown to impair reversal learning (Cools et al., 2001), possibly due to abnormal reward-related processing in the ventral striatum (Cools et al., 2007). Interestingly, administration of the DAT blocker methylphenidate resulted in similar impairment in healthy volunteers depending on the degree to which the drug increased dopamine release (Clatworthy et al., 2009). Thus, several lines of functional evidence have associated higher levels of dopamine with increased reward sensitivity and decreased behavioral flexibility. This would concur with our finding if the 9R allele is accompanied by higher dopamine activity or sensitivity, which is in line with studies showing enhanced reward-related responses in the ventromedial striatum in 9R carriers (Aarts et al., 2010; Dreher et al., 2009).

Finally, a second puzzle with respect to the effects of the *DAT1* genotype was that the relationship between performance during acquisition and perseveration during reversals actually reverses sign as a function of genotype. The computational model explains this as the tradeoff between two opposing effects. For low ρ , as observed in 10R homozygotes, the computational model approaches standard temporal difference learning. In such a model, as discussed above, performance on the acquisition and reversal phases are coupled by a common, local adjustment mechanism, with the degree of correct choices versus errors in both phases determined by choice randomness (the inverse temperature) and the sluggishness of adjustment (the learning rate). This produces a negative correlation between correct choices at acquisition and errors at reversal (equivalently, a positive correlation between errors in either phase). However, as

described above, for high ρ , as in the 9R genotype, the experience weighting mechanism produces the opposite effect. That is, increased choice of the correct stimulus during acquisition will lead to increased perseveration on reversal and therefore predict a positive relationship between the two.

Conclusions

This study revealed a functional double dissociation between the effects of polymorphisms in regulatory regions of the *SERT* and *DAT1* genes. We showed that within the same individuals, *SERT* is involved in behavioral adaptation following losses, whereas *DAT1* plays a role in experience-based perseveration. Our results provide strong and direct evidence for a suggested, but hitherto untested, functional dissociation, but fail to find a direct opponency between serotonin and dopamine systems.

EXPERIMENTAL PROCEDURES

Subjects

This study was part of the Brain Imaging Genetics (BIG) project at the Donders Institute for Brain, Cognition and Behavior, Radboud University Nijmegen. In the current study, 810 healthy, predominantly right-handed, Caucasian, highly educated subjects completed an online probabilistic reversal learning task among a set of other tests (60.4% female; age 26.3 ± 11.1 years (mean \pm SD); see [Table S1](#) for full demographic information). The study was approved by the local ethics committee (CMO 2001/095) and written informed consent was obtained from all subjects prior to participation.

Probabilistic Reversal Learning Task

Visual stimuli were probabilistically associated with positive (green, happy emoticon) and negative (red, sad/angry emoticon) feedback ([Figure 1](#)). We will refer to these positive and negative feedback events as “reward” and “punishment,” consistent with prior literature and the psychological definition of their tendency to increase/reduce response tendencies. On each trial, two stimuli were presented in two of four locations (left, right, top, or bottom of screen) and the subject was asked to select the usually rewarded stimulus with a mouse click. Choosing the correct stimulus (defined as the stimulus chosen on the first trial) resulted in a 70:30 ratio of reward/punishment. The incorrect stimulus resulted in the reverse (30:70) ratio. Thus, on 30% of trials subjects received “misleading” feedback. After 40 trials the reinforcement contingencies reversed, so that the frequently rewarded stimulus now became frequently punished and vice versa. Each subject completed a pseudorandom fixed sequence of 80 trials. Subjects were instructed that the identity of the correct stimulus could change, but received no information as to how often such a change might occur (for details see [Supplemental Experimental Procedures](#)).

Genotyping

Details of DNA extraction from the saliva samples and genotyping are described in the [Supplemental Experimental Procedures](#). For *DAT1*, two alleles of interest were analyzed: the common 10R allele and the rarer 9R allele. The insertion/deletion polymorphism in the *SERT* promoter region (5HTTLPR) was genotyped for the long (S) or short (L) alleles in combination with the single nucleotide polymorphism rs25531 A/G substitution in the same region. For the behavioral analysis, we used a biallelic model, where the S allele was grouped with the rare L_G allele (indicated as S'), given that the G-substitution in the L allele results in reduced expression more similar to the S allele ([Hu et al., 2006](#); [Praschak-Rieder et al., 2007](#)). L_A alleles were indicated as L'. Given the large sample size, all genotypes could be analyzed separately, which enabled testing for dose-dependent gene effects.

Behavioral Data Analysis

In all analyses, sex, age, and education level were included as covariates of no interest. The statistical significance threshold for all tests was $p = 0.05$, using a

Bonferroni correction where appropriate. To increase sensitivity, we did not use a Bonferroni correction for any of the control analyses.

Using the χ^2 test, we assessed whether there were any differences between genotype groups in the proportion of subjects passing the acquisition learning criterion of eight consecutive correct responses, which we report in the [Supplemental Experimental Procedures](#), where we also report baseline effects of task engagement/learning for both the pass and fail groups.

Behavioral Measures

Effects of reinforcement on subsequent choice were operationalized as the probability of repeating responses after reward (“win-stay”) and shifting responses after punishment (“lose-shift”) ([Figure 1A](#)).

Errors during the reversal phase were divided into two types. Perseverative errors were defined as two or more consecutive incorrect choices of the previously rewarded stimulus. Thus, perseverative errors required subjects to erroneously stay with the previously correct stimulus, despite punishment. The remaining errors during the reversal phase were defined as “chance errors.” The number of correct choices during acquisition was used as a measure reflecting the reinforcement history and value of the now incorrect stimulus at the start of reversal, in other words, it reflected how “stamped in” the choice of the initially correct stimulus was.

In a first analysis, win-stay, lose-shift, and perseveration rates were mean-corrected and entered in a repeated-measures GLM to assess any differential effects of the polymorphisms on these three measures. Learning criterion attainment and gender were included as fixed factors of no interest, and age and education were included as covariates of no-interest. The Huyn-Feldt correction was used when significant nonsphericity was detected. After significant interactions of *SERT* and *DAT1* with these behavioral measures were established, further analyses were used to determine the nature of these effects.

Win-Stay/Lose-Shift

Lose-shift and win-stay rates were entered as dependent variables in univariate ANOVA, with genotype for each polymorphism, learning criterion attainment (supplement) and gender as fixed effects, including all pairwise interactions. Age and education were included as covariates of no-interest. For significant effects ($p < 0.05$) post hoc pairwise t tests of the different genotypes were conducted to establish the nature of the genotype effects. Again, for significant effects, we then assessed the specificity with respect to the phase of the experiment (acquisition versus reversal) and the feedback validity, in two mixed repeated-measures ANOVAs with the same factors (*DAT1*, *SERT*, learning criterion attainment). For feedback validity, trials were divided into valid trials (win on a correct response, or loss on an incorrect response) and invalid trials. For task phase, trials were divided into acquisition and reversal phases of the task. Due to the small total number of trials it was not possible to perform this analysis in a single 2×2 factorial analysis.

Perseveration

The effect of genotype on the perseverative error rate was assessed using a hierarchical regression analysis with three sets of regressors: (1) regressors of no interest: sex, age, and education; (2) main effects: *DAT1* and *SERT* genotype and acquisition score; and (3) interactions: *DAT1* \times acquisition score, *SERT* \times acquisition score, and *DAT1* \times *SERT*. The same analysis was repeated for chance errors to establish the selectivity of the effect. We confirmed any gene-dose effects using a robust regression on the perseverative error rates versus acquisition scores for each genotype (Cauchy weighting, implemented in MATLAB 2011A). To ascertain that any observed effects on perseveration could not be explained by differences in acquisition, we assessed genotype effects on two basic measures of learning: (1) proportion of subjects passing a strict learning criterion of eight consecutive correct responses, using a χ^2 test, and (2) acquisition score, using an ANOVA.

Model-Based Analysis

Experience-Weighted Attraction Model

To understand the effects of *DAT1* on perseveration in the context of reinforcement learning, we used an augmented version of a standard Rescorla-Wagner model of learning. The key feature of this model is learning that is weighted by an experience weight. In this model, perseveration on reversal could occur because of an increasing reluctance to update the value of stimuli/choices every time they are chosen. Simplified to remove features unrelated to the

present study, the experience-weighted attraction (EWA) model of [Camerer and Ho \(1999\)](#) is described by the following equations:

$$n_{c,t} = n_{c,t-1} \times \rho + 1, \quad (\text{Equation 1})$$

and

$$V_{c,t} = (V_{c,t-1} \times \varphi \times n_{c,t-1} + \lambda_{t-1}) / n_{c,t}. \quad (\text{Equation 2})$$

Here, $n_{s,t}$ is the “experience weight” of stimulus s (blue or yellow) on trial t , which is updated on every trial, using the experience decay factor ρ . $V_{c,t}$ is the value of choice c on trial t , $\lambda_t \in \{0, 1\}$ for the outcome received in response to that choice and φ is the decay factor for the previous payoffs, equivalent to the learning rate in the Rescorla-Wagner model. In particular, note that for $\rho = 0$, $n_{c,t}$ is everywhere 1, and the model reduces to Rescorla-Wagner. For $\rho > 0$, the experience weights promote more sluggish updating with time. Note that a rearrangement of the parameters is required to see the equivalence between these equations and Rescorla-Wagner. The Rescorla-Wagner learning rate, usually denoted α , is here equivalent to $(1 - \varphi)$. Moreover, the softmax inverse temperature β , below, is equivalent to the product $\beta\alpha$ in Rescorla-Wagner. This is because the values $V_{c,t}$ learned here are scaled by a constant factor of $1/\alpha$ relative to those learned by their Rescorla-Wagner equivalents. This rescaling makes the model more numerically stable at small α .

RP Model

The hypothesis reflected by this model is that perseverative behavior is caused by reduced learning from punishment, where punishment to the previously rewarded stimulus has little effect, resulting in a failure to devalue this stimulus. This model is described by the following equations:

$$V_{c,t} = V_{c,t-1} + \alpha^{\text{pun}} \times (\lambda_{t-1} - V_{c,t-1}) + \alpha^{\text{rew}} \times (\lambda_{t-1} - V_{c,t-1}) \quad (\text{Equation 3})$$

and

$$V_{-c,t} = V_{-c,t-1}, \quad (\text{Equation 4})$$

where α^{pun} is the punishment learning rate (0 on reward trials), and α^{rew} is the learning rate for reward (0 on punishment trials). $V_{-c,t}$ is the value of the unchosen option. Note that only the chosen stimulus is updated.

Action Selection

For both models, to select an action based on the computed values, we used a softmax choice function to compute the probability of each choice. For a given set of parameters, this equation allows us to compute the probability of the next choice being “ i ” given the previous choices:

$$p(C_{t+1} = i) = \frac{e^{\beta Q(c=i,t+1)}}{\sum_j e^{\beta Q(c=j,t+1)}}. \quad (\text{Equation 5})$$

Here, β is the inverse temperature parameter.

Model Fitting

For both models, we fit all parameters separately to the choices of each individual ([RP: $\alpha_{\text{pun}}, \alpha_{\text{rew}}, \beta$; EWA: ϕ, ρ, β]). To facilitate stable estimation across so large a group of subjects, we used weakly informative priors ([Table 1](#)) to regularize the estimated priors toward realistic ones. Thus we use maximum a posteriori (MAP; rather than maximum likelihood) estimation ([Daw, 2011](#)). In particular, we optimized model parameters by minimizing the negative log posterior of the observed choice sequence, given the previously observed outcomes, with respect to different settings of the model parameters.

Model Comparison

To investigate which model best described the data, we computed the Bayesian evidence E_m or probability of the model given the data for each model, using the Laplace approximation ([Kass and Raftery, 1995](#)):

$$E_m \approx \log p(\hat{\theta}_m) + \log p(c_{1:T} | \hat{\theta}_m) + \frac{1}{2} G_m \log 2\pi - \frac{1}{2} \log |H_m|. \quad (\text{Equation 6})$$

This quantity, like the Bayesian Information Criterion ([Schwarz, 1978](#)), which can be derived from it via a further approximation) scores each model accord-

ing to its fit to the data, penalized for overfitting due to optimizing the models’ parameters. Here, $\hat{\theta}_m$ are the best fitting MAP parameters, $p(\hat{\theta}_m)$ is the value of the prior on the MAP parameters, $p(c_{1:T} | \hat{\theta}_m)$ is the likelihood of the series of observed choices on trials 1- T , G_m is the number of parameters in the model m , and $|H_m|$ is the determinant of the Hessian matrix of the second derivatives of the negative log posterior with respect to the parameters, evaluated at the MAP estimate.

This Bayesian evidence can then be used to compare models of different complexity by correctly penalizing models for their differing (effective) number of free parameters. Having computed this score separately for each subject and model, to compare the fits at the population level, we used the random-effects Bayesian model selection procedure ([Stephan et al., 2009](#)), in which model identity is taken as a random effect—i.e., each subject might instantiate a different model—and the relative proportions of each model across the population are estimated. From these, we derive the exceedance probability XP_m , i.e., the posterior probability, given the data, that a particular model m is the most common model in the group.

Significance Tests on Estimated Model Parameters

To assess evidence for dose-dependent effects of the *DAT1* polymorphism on any of the model parameters of the best-fitting model, we used Jonckheere-Terpstra for ordered alternatives, a nonparametric test due to non-Gaussianity of the parameters. Significance is reported at a very strict Bonferroni-corrected significance level of 0.0083 (2 genes \times 3 parameters). For completeness, we also tested whether fitted parameter values in the losing model differed with *DAT1* genotype.

Model Simulations

To assess whether the model could replicate the behavioral findings, we generated trial-by-trial choices using the fitted parameters of the best fitting model. We then analyzed these choices in the same way as the original data, again using robust regression analyses.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, two figures, and three tables and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2013.08.030>.

ACKNOWLEDGMENTS

We thank Sabine Kooijman for logistic support; Angelien Heister, Remco Makkinje, and Marlies Naber for genotyping; and Bradley Doll, Sean Fallon, Michael Frank, Guillaume Sescousse, and Jennifer Cook for insightful discussions and feedback. This work makes use of the Brain Imaging Genetics (BIG) database, first established in Nijmegen, the Netherlands, in 2007. This resource is now part of Cognomics (<http://www.cognomics.nl>), a joint initiative by researchers of the Donders Centre for Cognitive Neuroimaging, the Human Genetics and Cognitive Neuroscience departments of the Radboud University Medical Centre and the Max Planck Institute for Psycholinguistics in Nijmegen. The Cognomics Initiative is supported by the participating departments and centres and by external grants: the Biobanking and Biomolecular Resources Research Infrastructure (Netherlands) (BBMRI-NL), the Hersenstichting Nederland, and the Netherlands Organisation for Scientific Research. This study was also supported by a Research Vidi Grant to R.C. and a Research Veni Grant to H.d.O. from the Innovational Research Incentives Scheme of the Netherlands Organisation for Scientific Research as well as a Human Frontiers Science Program grant to Kae Nakamura, N.D., and R.C., and a James McDonnell scholar award to both R.C. and N.D. We wish to thank all who kindly participated in this research.

Accepted: August 26, 2013

Published: November 20, 2013

REFERENCES

Aarts, E., Roelofs, A., Franke, B., Rijpkema, M., Fernández, G., Helmich, R.C., and Cools, R. (2010). Striatal dopamine mediates the interface between

- motivational and cognitive control in humans: evidence from genetic imaging. *Neuropsychopharmacology* 35, 1943–1951.
- Balleine, B.W., and O'Doherty, J.P. (2010). Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology* 35, 48–69.
- Bari, A., Theobald, D.E., Caprioli, D., Mar, A.C., Aidoo-Micah, A., Dalley, J.W., and Robbins, T.W. (2010). Serotonin modulates sensitivity to reward and negative feedback in a probabilistic reversal learning task in rats. *Neuropsychopharmacology* 35, 1290–1301.
- Boulougouris, V., and Robbins, T.W. (2010). Enhancement of spatial reversal learning by 5-HT_{2C} receptor antagonism is neuroanatomically specific. *J. Neurosci.* 30, 930–938.
- Boulougouris, V., Glennon, J.C., and Robbins, T.W. (2008). Dissociable effects of selective 5-HT_{2A} and 5-HT_{2C} receptor antagonists on serial spatial reversal learning in rats. *Neuropsychopharmacology* 33, 2007–2019.
- Boulougouris, V., Castañé, A., and Robbins, T.W. (2009). Dopamine D₂/D₃ receptor agonist quinpirole impairs spatial reversal learning in rats: investigation of D₃ receptor involvement in persistent behavior. *Psychopharmacology (Berl.)* 202, 611–620.
- Boureau, Y.L., and Dayan, P. (2011). Opponency revisited: competition and cooperation between dopamine and serotonin. *Neuropsychopharmacology* 36, 74–97.
- Camerer, C., and Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica* 67, 827–874.
- Caspi, A., Hariri, A.R., Holmes, A., Uher, R., and Moffitt, T.E. (2010). Genetic sensitivity to the environment: the case of the serotonin transporter gene and its implications for studying complex diseases and traits. *Am. J. Psychiatry* 167, 509–527.
- Chamberlain, S.R., Müller, U., Blackwell, A.D., Clark, L., Robbins, T.W., and Sahakian, B.J. (2006). Neurochemical modulation of response inhibition and probabilistic learning in humans. *Science* 311, 861–863.
- Clarke, H.F., Walker, S.C., Dalley, J.W., Robbins, T.W., and Roberts, A.C. (2007). Cognitive inflexibility after prefrontal serotonin depletion is behaviorally and neurochemically specific. *Cereb. Cortex* 17, 18–27.
- Clatworthy, P.L., Lewis, S.J., Brichard, L., Hong, Y.T., Izquierdo, D., Clark, L., Cools, R., Aigbirhio, F.I., Baron, J.C., Fryer, T.D., and Robbins, T.W. (2009). Dopamine release in dissociable striatal subregions predicts the different effects of oral methylphenidate on reversal learning and spatial working memory. *J. Neurosci.* 29, 4690–4696.
- Collins, A.G., and Frank, M.J. (2012). How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *Eur. J. Neurosci.* 35, 1024–1035.
- Cools, R., Barker, R.A., Sahakian, B.J., and Robbins, T.W. (2001). Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cereb. Cortex* 11, 1136–1143.
- Cools, R., Altamirano, L., and D'Esposito, M. (2006). Reversal learning in Parkinson's disease depends on medication status and outcome valence. *Neuropsychologia* 44, 1663–1673.
- Cools, R., Lewis, S.J.G., Clark, L., Barker, R.A., and Robbins, T.W. (2007). L-DOPA disrupts activity in the nucleus accumbens during reversal learning in Parkinson's disease. *Neuropsychopharmacology* 32, 180–189.
- Cools, R., Roberts, A.C., and Robbins, T.W. (2008a). Serotonergic regulation of emotional and behavioural control processes. *Trends Cogn. Sci.* 12, 31–40.
- Cools, R., Robinson, O.J., and Sahakian, B. (2008b). Acute tryptophan depletion in healthy volunteers enhances punishment prediction but does not affect reward prediction. *Neuropsychopharmacology* 33, 2291–2299.
- Cools, R., Frank, M.J., Gibbs, S.E., Miyakawa, A., Jagust, W., and D'Esposito, M. (2009). Striatal dopamine predicts outcome-specific reversal learning and its sensitivity to dopaminergic drug administration. *J. Neurosci.* 29, 1538–1543.
- Cools, R., Nakamura, K., and Daw, N.D. (2011). Serotonin and dopamine: unifying affective, motivational, and decision functions. *Neuropsychopharmacology* 36, 98–113.
- Costa, A., Riedel, M., Müller, U., Möller, H.J., and Ettinger, U. (2011). Relationship between SLC6A3 genotype and striatal dopamine transporter availability: a meta-analysis of human single photon emission computed tomography studies. *Synapse* 65, 998–1005.
- Crockett, M.J., Clark, L., and Robbins, T.W. (2009). Reconciling the role of serotonin in behavioral inhibition and aversion: acute tryptophan depletion abolishes punishment-induced inhibition in humans. *J. Neurosci.* 29, 11993–11999.
- Daw, N.D., Kakade, S., and Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural Netw.* 15, 603–616.
- Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8, 1704–1711.
- Daw, N.D. (2011). Trial by trial data analysis using computational models. In *Decision Making, Affect, and Learning: Attention and Performance XXIII*, M.R. Delgado, E.A. Phelps, and T.W. Robbins, eds. (Oxford: Oxford University Press), pp. 3–48.
- Dayan, P., and Huys, Q.J. (2008). Serotonin, inhibition, and negative mood. *PLoS Comput. Biol.* 4, e4.
- Dayan, P., Kakade, S., and Montague, P.R. (2000). Learning and selective attention. *Nat. Neurosci. Suppl.* 3, 1218–1223.
- Deakin, J.F.W., and Graeff, F.G. (1991). 5-HT and mechanisms of defence. *J. Psychopharmacol. (Oxford)* 5, 305–315.
- Dodds, C.M., Müller, U., Clark, L., van Loon, A., Cools, R., and Robbins, T.W. (2008). Methylphenidate has differential effects on blood oxygenation level-dependent signal related to cognitive subprocesses of reversal learning. *J. Neurosci.* 28, 5976–5982.
- Dreher, J.C., Kohn, P., Kolachana, B., Weinberger, D.R., and Berman, K.F. (2009). Variation in dopamine genes influences responsiveness of the human reward system. *Proc. Natl. Acad. Sci. USA* 106, 617–622.
- Everitt, B.J., and Robbins, T.W. (2005). Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nat. Neurosci.* 8, 1481–1489.
- Fiorillo, C.D., Tobler, P.N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299, 1898–1902.
- Forbes, E.E., Brown, S.M., Kimak, M., Ferrell, R.E., Manuck, S.B., and Hariri, A.R. (2009). Genetic variation in components of dopamine neurotransmission impacts ventral striatal reactivity associated with impulsivity. *Mol. Psychiatry* 14, 60–70.
- Frank, M.J., Seeberger, L.C., and O'Reilly, R. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science* 306, 1940–1943.
- Frank, M.J., Moustafa, A.A., Haughey, H.M., Curran, T., and Hutchison, K.E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proc. Natl. Acad. Sci. USA* 104, 16311–16316.
- Franke, B., Vasquez, A.A., Johansson, S., Hoogman, M., Romanos, J., Boreatti-Hümmer, A., Heine, M., Jacob, C.P., Lesch, K.P., Casas, M., et al. (2010). Multicenter analysis of the SLC6A3/DAT1 VNTR haplotype in persistent ADHD suggests differential involvement of the gene in childhood and persistent ADHD. *Neuropsychopharmacology* 35, 656–664.
- Gizer, I.R., Ficks, C., and Waldman, I.D. (2009). Candidate gene studies of ADHD: a meta-analytic review. *Hum. Genet.* 126, 51–90.
- Goto, Y., and Grace, A.A. (2005). Dopaminergic modulation of limbic and cortical drive of nucleus accumbens in goal-directed behavior. *Nat. Neurosci.* 8, 805–812.
- Hariri, A.R., and Holmes, A. (2006). Genetics of emotional regulation: the role of the serotonin transporter in neural function. *Trends Cogn. Sci.* 10, 182–191.
- Houk, J.C., Adams, J.L., and Barto, A.G. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. In *Models of Information Processing in the Basal Ganglia*, J.C. Houk, J.L. Davis, and D.G. Beiser, eds. (Cambridge, MA: MIT Press), pp. 249–270.
- Hu, X.Z., Lipsky, R.H., Zhu, G., Akhtar, L.A., Taubman, J., Greenberg, B.D., Xu, K., Arnold, P.D., Richter, M.A., Kennedy, J.L., et al. (2006). Serotonin

- transporter promoter gain-of-function genotypes are linked to obsessive-compulsive disorder. *Am. J. Hum. Genet.* **78**, 815–826.
- Izquierdo, A., Newman, T.K., Higley, J.D., and Murray, E.A. (2007). Genetic modulation of cognitive flexibility and socioemotional behavior in rhesus monkeys. *Proc. Natl. Acad. Sci. USA* **104**, 14128–14133.
- Jedema, H.P., Gianaros, P.J., Greer, P.J., Kerr, D.D., Liu, S., Higley, J.D., Suomi, S.J., Olsen, A.S., Porter, J.N., Lopresti, B.J., et al. (2010). Cognitive impact of genetic variation of the serotonin transporter in primates is associated with differences in brain morphology rather than serotonin neurotransmission. *Mol. Psychiatry* **15**, 512–522.
- Kass, R.E., and Raftery, A.E. (1995). Bayes Factors. *J. Am. Stat. Assoc.* **90**, 773–795.
- Lawrence, A.D., Sahakian, B.J., Rogers, R.D., Hodge, J.R., and Robbins, T.W. (1999). Discrimination, reversal, and shift learning in Huntington's disease: mechanisms of impaired response selection. *Neuropsychologia* **37**, 1359–1374.
- Lesch, K.P., Bengel, D., Heils, A., Sabol, S.Z., Greenberg, B.D., Petri, S., Benjamin, J., Müller, C.R., Hamer, D.H., and Murphy, D.L. (1996). Association of anxiety-related traits with a polymorphism in the serotonin transporter gene regulatory region. *Science* **274**, 1527–1531.
- Maia, T.V., and Frank, M.J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nat. Neurosci.* **14**, 154–162.
- Mill, J., Asherson, P., Browes, C., D'Souza, U., and Craig, I. (2002). Expression of the dopamine transporter gene is regulated by the 3' UTR VNTR: evidence from brain and lymphocytes using quantitative RT-PCR. *Am. J. Med. Genet.* **114**, 975–979.
- Montague, P.R., Dayan, P., and Sejnowski, T.J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* **16**, 1936–1947.
- Murphy, F.C., Michael, A., Robbins, T.W., and Sahakian, B.J. (2003). Neuropsychological impairment in patients with major depressive disorder: the effects of feedback on task performance. *Psychol. Med.* **33**, 455–467.
- Palminteri, S., Clair, A.H., Mallet, L., and Pessiglione, M. (2012). Similar improvement of reward and punishment learning by serotonin reuptake inhibitors in obsessive-compulsive disorder. *Biol. Psychiatry* **72**, 244–250.
- Parkinson, J.A., Olmstead, M.C., Burns, L.H., Robbins, T.W., and Everitt, B.J. (1999). Dissociation in effects of lesions of the nucleus accumbens core and shell on appetitive pavlovian approach behavior and the potentiation of conditioned reinforcement and locomotor activity by D-amphetamine. *J. Neurosci.* **19**, 2401–2411.
- Praschak-Rieder, N., Kennedy, J., Wilson, A.A., Hussey, D., Boovariwala, A., Willeit, M., Ginovart, N., Tharmalingam, S., Masellis, M., Houle, S., and Meyer, J.H. (2007). Novel 5-HTTLPR allele associates with higher serotonin transporter binding in putamen: a [¹¹C] DASB positron emission tomography study. *Biol. Psychiatry* **62**, 327–331.
- Robinson, O.J., Standing, H.R., DeVito, E.E., Cools, R., and Sahakian, B.J. (2010). Dopamine precursor depletion improves punishment prediction during reversal learning in healthy females but not males. *Psychopharmacology (Berl.)* **211**, 187–195.
- Roiser, J.P., de Martino, B., Tan, G.C., Kumaran, D., Seymour, B., Wood, N.W., and Dolan, R.J. (2009). A genetically mediated bias in decision making driven by failure of amygdala control. *J. Neurosci.* **29**, 5985–5991.
- Rutledge, R.B., Lazzaro, S.C., Lau, B., Myers, C.E., Gluck, M.A., and Glimcher, P.W. (2009). Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. *J. Neurosci.* **29**, 15104–15114.
- Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* **275**, 1593–1599.
- Schwarz, G. (1978). Estimating dimensions of a model. *Ann. Stat.* **6**, 461–464.
- Soubrie, P. (1986). Reconciling the role of central serotonin neurons in human and animal behavior. *Behav. Brain Sci.* **9**, 319–335.
- Spencer, T.J., Biederman, J., Faraone, S.V., Madras, B.K., Bonab, A.A., Dougherty, D.D., Batchelder, H., Clarke, A., and Fischman, A.J. (2013). Functional genomics of attention-deficit/hyperactivity disorder (ADHD) risk alleles on dopamine transporter binding in ADHD and healthy control subjects. *Biol. Psychiatry* **74**, 84–89.
- Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., and Friston, K.J. (2009). Bayesian model selection for group studies. *Neuroimage* **46**, 1004–1017.
- Suri, R.E., and Schultz, W. (1999). A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience* **91**, 871–890.
- Thorndike, E. (1911). *Animal Intelligence: Experimental Studies*. (New York: Macmillan).
- Vallender, E.J., Lynch, L., Novak, M.A., and Miller, G.M. (2009). Polymorphisms in the 3' UTR of the serotonin transporter are associated with cognitive flexibility in rhesus macaques. *Am. J. Med. Genet. B. Neuropsychiatr. Genet.* **150B**, 467–475.
- van de Giessen, E., de Win, M.M., Tanck, M.W., van den Brink, W., Baas, F., and Booij, J. (2009). Striatal dopamine transporter availability associated with polymorphisms in the dopamine transporter gene SLC6A3. *J. Nucl. Med.* **50**, 45–52.
- van Dyck, C.H., Malison, R.T., Jacobsen, L.K., Seibyl, J.P., Staley, J.K., Laruelle, M., Baldwin, R.M., Innis, R.B., and Gelernter, J. (2005). Increased dopamine transporter availability associated with the 9-repeat allele of the SLC6A3 gene. *J. Nucl. Med.* **46**, 745–751.
- Walker, S.C., Robbins, T.W., and Roberts, A.C. (2009). Differential contributions of dopamine and serotonin to orbitofrontal cortex function in the marmoset. *Cereb. Cortex* **19**, 889–898.
- Willeit, M., and Praschak-Rieder, N. (2010). Imaging the effects of genetic polymorphisms on radioligand binding in the living human brain: a review on genetic neuroreceptor imaging of monoaminergic systems in psychiatry. *Neuroimage* **53**, 878–892.