

Human substantia nigra and ventral tegmental area involvement in computing social error signals during the ultimatum game

Sébastien Héту¹, Yi Luo¹, Kimberlee D'Ardenne¹, Terry Lohrenz¹, and P. Read Montague^{1,2}

¹Virginia Tech Carilion Research Institute, 2 Riverside Circle, Roanoke, VA 24016, USA and ²Wellcome Trust Centre for Neuroimaging, University College London, 12 Queen Square, London WC1N 3BG, UK

Correspondence should be addressed to P. Read Montague, Virginia Tech Carilion Research Institute, 2 Riverside Circle, Roanoke, VA 24016, USA and Wellcome Trust Centre for Neuroimaging, University College London, 12 Queen Square, London WC1N 3BG, UK. E-mail: read@vtc.vt.edu

Abstract

As models of shared expectations, social norms play an essential role in our societies. Since our social environment is changing constantly, our internal models of it also need to change. In humans, there is mounting evidence that neural structures such as the insula and the ventral striatum are involved in detecting norm violation and updating internal models. However, because of methodological challenges, little is known about the possible involvement of midbrain structures in detecting norm violation and updating internal models of our norms. Here, we used high-resolution cardiac-gated functional magnetic resonance imaging and a norm adaptation paradigm in healthy adults to investigate the role of the substantia nigra/ventral tegmental area (SN/VTA) complex in tracking signals related to norm violation that can be used to update internal norms. We show that the SN/VTA codes for the norm's variance prediction error (PE) and norm PE with spatially distinct regions coding for negative and positive norm PE. These results point to a common role played by the SN/VTA complex in supporting both simple reward-based and social decision making.

Key words: decision making; fMRI; prediction error; social neuroscience; ultimatum game

Introduction

Social norms represent the rules of behavior normally associated with certain social situations: they are the behavioral rules that one *ought* to conform to (Hechter and Opp, 2001). They thus combine empirical (we think that most people conform) and normative (we think that people ought to conform) expectations. Social norms play an essential role in our societies as they are models of shared expectations we can use to summarize the social environment and guide our decision process. Since societies are constantly changing, our models also need to change. This implies that the human brain can (1) compute a shared norm about what is expected, (2) detect

deviations from that norm and (3) choose the best actions to correct these deviations (Montague and Lohrenz, 2007). These corrective actions can be to (1) modify the environment or (2) change their internal norm in order to bring it closer to the present state of their environment.

Specific brain regions track values important for the computation needed to adapt ones internal norms to a changing environment. In addition to playing a role in detecting norm deviations (Chang and Sanfey, 2013), the insula has been shown to play a crucial role in norm updating as people with insula damage are slower to adapt to changes in norm than healthy controls (Gu et al., 2015). The ventromedial prefrontal/medial

Received: 26 September 2016; Revised: 24 July 2017; Accepted: 7 August 2017

© The Author (2017). Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

orbitofrontal cortex (mOFC) and ventral striatum have also been shown to be implicated in norm adaptation by encoding differences between expected norms and new information about the state of the environment (Xiang et al., 2013). These differences/deviations can be considered as norm prediction errors (PEs). Norm PE signals give the direction and an estimate of the amount by which we have to move our internal norm in order to adapt to the environment. Interestingly, positive norm PE and negative norm PE (respectively when the world is better or worse than we expected) seem to be, at least partially, computed in spatially distinct regions (Xiang et al., 2013).

A large body of work has shown that the brain keeps track of the uncertainty about the values it computes (reviewed in Vilares and Kording, 2011; Bach and Dolan, 2012; Ma and Jazayeri, 2014). Uncertainty is often operationalized by variance or risk measures which have been shown to be represented by the activity of dopaminergic regions (Fiorillo et al., 2003; Preusschoff et al., 2006) and the anterior insula (Preusschoff et al., 2008). Furthermore, an error signal related to risk updating, the risk PE was shown to be tracked by the anterior insula (Preusschoff et al., 2008; d'Acremont et al., 2009). The role of the anterior insula and ventral striatum in updating estimates about uncertainty has been extended to the social domain by Xiang et al. (2013) who showed that the activity within these regions covaried with variance PE signal related to social norms.

In concept, norm and variance PEs are strikingly similar to reward PEs that play an important role in reinforcement learning (Montague et al., 1996; Schultz et al., 1997) and both norm (Xiang et al., 2013) and reward (e.g. McClure et al., 2003; O'Doherty et al., 2003; Rolls et al., 2008) PE have been shown to be encoded in similar regions targeted by the dopamine system. Blood-oxygen-level dependent (BOLD) signals correlating with reward PEs have also been reported within the human midbrain's substantia nigra/ventral tegmental area (SN/VTA) complex (Dreher et al., 2006; D'Ardenne et al., 2008, 2013) which is rich in dopamine neurons. Importantly, studies that found signals related to reward PE in the SN/VTA have only used non-social tasks where the participants were not interacting with other individuals. There is evidence that computing information in a social context can be different than in a non-social context. Indeed, in social tasks involving monetary 'rewards' such as the ultimatum game (UG) and the prisoner's dilemma game (see Camerer, 2003 for details about these tasks), participants behave differently (Rilling et al., 2004), and have specific physiological (Van't Wout et al., 2006) and neural responses (Rilling et al., 2004) if they think they are playing a human vs. a computer player. This suggests that human participants process/represent monetary 'rewards' differently when in a social context. To date no study has specifically looked at the role of the midbrain in social interactions and particularly if it is involved in social norm adaptation by computing PE or variance PE signals related to norm violations (but see Klucharev et al., 2009 for a study on social conformity which used an imaging approach not tailored for the midbrain). Indeed, functional magnetic resonance imaging (fMRI) of the human midbrain is complicated by methodological challenges (reviewed in Beissner, 2015) related to the small size of these structures (Eapen et al., 2011; Lawson et al., 2013) and their proximity to important blood vessels which may affect imaging signals (Enzmann and Pelc, 1992; Dagli et al., 1999).

Here, we had healthy individuals play the role of a Responder in an UG task in which the offers were manipulated to induce changes in their expectations (norms). By recording their brain activity using high-resolution fMRI of the midbrain,

we were able to study the role of the SN/VTA in the computations underlying social norm adaptation.

Materials and methods

Participants

Fifty participants (30 females) were recruited in the Roanoke (VA) region. The average age of the participants was 27.5 (s.d. = 10.3; range 19–58). The study was approved by the Virginia Tech Institutional Review Board and written consent was given by all participants.

Task procedure

Participants completed a modified version of the UG. In the original UG, a Proposer gets a 20\$ initial endowment, decides on a split proportion to send to a Responder who in turn has to choose to accept or reject the proposed split. If the Responder accepts, each gets the proposed amount, if the Responder rejects both get 0\$. Participants were instructed that they would play the role of the Responder in 60 such exchanges and that they would play with a different randomly selected Proposer each time. Unknown to the participants, the offers they received were not from human Proposers but rather generated by a computer algorithm in order to condition/manipulate the Responder's expectations (norms). We used three offer distributions, each truncated Gaussians (1) high or 'generous' offers¹ (mean = 12\$, s.d. = 1.5\$); (2) low or 'greedy' offers (mean = 4\$, s.d. = 1.5\$); (3) medium or 'fair'² offers (mean = 8\$, s.d. = 1.5\$). Participants were randomly assigned to one of two groups. For the first 30 trials, the low-to-medium group (L2M) received offers taken from the low offers distribution whereas the high-to-medium group (H2M) received offers from the high offers distribution. In the last 30 trials, both groups received offers taken from the medium offers distribution creating a 'positive' contrast (i.e. expecting low offers but receiving medium offers) or a 'negative' (i.e. expecting high offers but receiving medium offers) contrast. We hypothesized that as in Xiang et al. (2013) participants previously conditioned on high vs low offers—experiencing a 'positive vs negative' contrast—would behave differently toward subsequent medium offers. Our choice of using a between subject design was driven by the fact that our paradigm is based on a conditioning effect. Hence, using a within subject design where participants would undergo both conditioning conditions would have potentially led to some 'contamination' of the second conditioning condition by the previous experience of high/low and medium offers.

1 Offers >50% (hyperfair) are rarer in the USA than unfair offers or equal splits (Henrich et al., 2006). However, they have been observed in various contexts (see the examples given in Hennig-Schmidt et al., 2008). We also want to point out that the hyperfair offers used in this study were taken from a distribution with a mean of 12 (not far from 50/50) and an s.d. of 1.5. Therefore, very unlikely hyperfair offers (e.g. 16–17–18–19\$) were very seldom presented. This was done in order to make our offers 'believable'. In fact, none of the participants mentioned after the task that they believed they were NOT playing with a human partner.

2 We propose that the use of the word *fair* to describe our distribution in the second half is more in line with our UG setup than *equal* as this distribution was not centered on 10\$ but rather 8\$. We decided to use a distribution with a mean lower than 10\$ because data have shown that Responders usually expect (what they think is fair) Proposers to give offers slightly lower than the 50–50 split (Henrich et al., 2006) suggesting that fairness ≠ equality (see also the paper by Starmans et al., 2017 on the differences between inequality and unfairness).

Participants were informed to treat each exchange equally as their payment would be based on the outcome of two randomly selected trials. The mean performance-based payment was 13.97\$ (s.d. = 5.47\$) plus the base payment of 20\$/h. For each trial, a preparation screen indicating the pairing to a new Proposer, an offer screen where the offer split was shown, a response screen where the participants were prompted to indicate their response (accept–reject), and a blank screen indicating the end of the trial were sequentially shown (Figure 1). At the end of 60% of the trials (every three of five trials), an emotion rating screen was presented and participants were prompted to rate their feelings about the offers they received using emoticons ranging from sad to happy on a 1–9 scale (adapted from Lang’s (1980) self-assessment manikin). Visual display of the task was back-projected onto a computer screen which participants viewed using a mirror placed inside the scanner. Choices were indicated by pressing a button on a hand-held response box. Stimuli presentation and participants’ behavioral responses collection were controlled by the NEMO software (Human Neuroimaging Laboratory, Virginia Tech Carilion Research Institute; <http://labs.vtc.vt.edu/hnl/software.html>).

Behavioral data analysis

Norm adaptation model. Based on data showing that healthy subjects have an internal representation of norms which can be modified when confronted to changes in their environment (Xiang et al., 2013; Gu et al., 2015) and that these norms can be used to guide choices (Chang and Sanfey, 2013; Sanfey et al., 2014), we assumed that the participants’ behavior could be modeled by their aversion to splits that deviate from their internal norm (see Supplementary material for details). This aversion is controlled by two parameters (Fehr and Schmidt, 1999; Bicchieri, 2005) that we estimated using their behaviors: α or ‘envy’ which represents the

unwillingness of participants to accept offers that are lower than their internal norm and β or ‘guilt’ which represents their unwillingness to accept offers that are too good (i.e. higher than their internal norm). We made the assumption that participants were Bayesian observers who modified their internal norm based on the offers they received throughout the task (see Supplementary material for details). This model has previously been used to model behavior in a similar norm adaptation task (Xiang et al., 2013). Using this model we estimated, for each participant on each trial, values related to social norms updating (Figure 2): the norm PE (difference between the offer and the participant’s social norm; $PE = \text{offer} - \text{norm}$); the positive norm PE [difference between the offer and the participant’s social norm when offers are bigger than the norm; $\text{posPE} = \max(\text{offer} - \text{norm}, 0)$], the negative norm PE [difference between the offer and the participant’s social norm when offers are smaller than the norm; $\text{negPE} = \min(\text{offer} - \text{norm}, 0)$]; and the norm’s variance/uncertainty PE (the difference between the observed social norms’ uncertainty and the participant’s expectation about this uncertainty—Preuschoff et al., 2008).

Norm adaptation behavioral analysis. As tested by the Kolmogorov–Smirnov test, the normality assumption of parametric tests was not met by our behavioral dataset (rejection rate and emotional rating). Therefore, we estimated the probability of observing a difference between our two experimental groups by using a bootstrapping method (Mooney et al., 1993). Differences were considered significant if the probability of obtaining the actual *t* statistics was <5% (two-tailed) along the permuted (10 000 iterations) distribution of *t* statistics (see Supplementary material for details).

fMRI acquisition and preprocessing

Imaging was done using an approach tailored for the study of the SN/VTA complex (D’Ardenne et al., 2008) (Figure 2). Images




Fig. 1. Description of the norm adjustment task. (A) On each trial of the UG, a Proposer has to decide how to divide a 20\$ endowment between himself and a Responder. The Responder can then accept or reject this split: if he accepts, the Proposer and the Responder receive their share of the split; if he rejects, both the Proposer and the Responder receive nothing. In our version of the UG, the participants played the role of the Responder. (B) Gaussian distributions from which the offers were sampled. Unknown to the participants, the Proposers’ offers were controlled by the experimenters. For the L2M, participants first received 30 offers sampled from an ‘greedy’ distribution (red; mean = 4\$, s.d. = 1.5\$) and then received 30 offers sampled from a ‘fair’ distribution (gray; mean = 8\$, s.d. = 1.5\$). For the H2M, participants first received 30 offers sampled from a ‘generous’ distribution (blue; mean = 12\$, s.d. = 1.5\$) and then received 30 offers sampled from a ‘fair’ distribution (gray). (C) Depiction of the task’s timeline and visual display. Each trial began with the presentation of a new partner (4 s). Offers were displayed for 4 s and then participants were prompted to respond (self-paced). On 60% of the trials, participants had to rate their feelings toward the offer from 1 (sad face) to 9 (happy face) at a self-paced speed. Intertrial intervals were between 2 and 4 s. L2M: low-to-medium group; H2M: high-to-medium group.

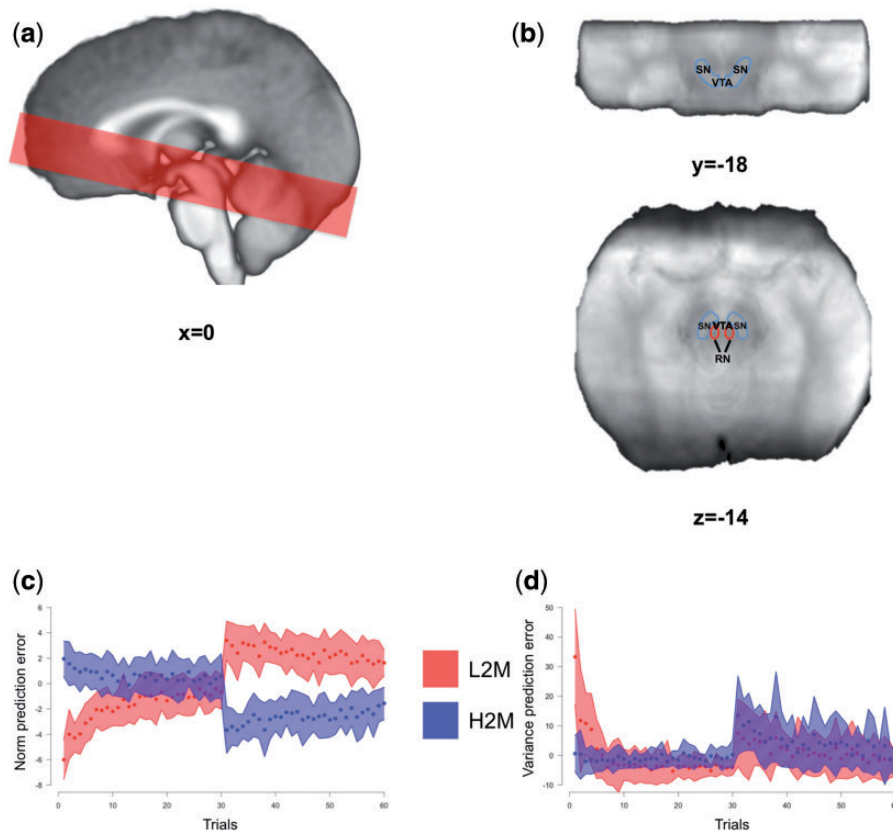


Fig. 2. Imaging details. (A) Example of the axial slab placement overlaid over the mean T1 image. (B) Identification of the SN/VTA complex in a coronal (top) and axial (bottom) view. The SN/VTA was drawn using the mean proton density MR image. (C, D) We looked for brain correlates of the norm PE (C) and risk/variance PE (D) signals. Mean values for the low-to-medium (red) and high-to-medium (blue) groups are depicted as dots and the shaded area represents \pm s.d. RN: red nucleus; SN: substantia nigra; VTA: ventral tegmental area; s.d.: standard deviation; L2M: low-to-medium group; H2M: high-to-medium group.

were acquired (see Supplementary material for details about acquisition and preprocessing) on a 3-Tesla TimTrio Siemens scanner at the Human Neuroimaging Laboratory-Virginia Tech Carilion Research Institute. High-resolution ($0.5 \times 0.5 \times 1$ mm) T1-weighted structural images were acquired using an MP-RAGE pulse sequence and functional images were acquired using a high-resolution (1.5×1.5 in plane $\times 1.9$ mm) cardiac-gated echoplanar imaging sequence. To facilitate the localization of the SN and VTA, a proton-density weighted image was also acquired. The borders of SN and VTA were visually identified on the group's mean proton-density image. Using the AFNI's drawing tool, an SN/VTA region of interest (ROI) was drawn. The resulting ROI comprised of 658 voxels (Figure 2). Preprocessing and fMRI analyses including despiking, slice-timing and motion corrections, smoothing and normalization were performed with AFNI (Cox, 1996).

After the preprocessing steps, 11 subjects were excluded from further analyses (2 because of acquisition problems, 6 had extensive head movements, 3 because of problems in the alignment step). This left 20 participants in the L2M group and 19 in the H2M group.

fMRI data analysis. General linear models (GLM) were specified for each subject. For all analyses, the design matrix included onsets of all visual stimuli (start screen, offer screen, response screen, emotion rating screen) and motor responses convolved with AFNI's gamma hemodynamic response function (event duration = 0 s). The six motion parameters and a vector of the TR lengths (which is unique for each subject) were also included

as covariates of non-interest. In order to assess if the brainstem was involved in the internal norm updating process, we looked for areas that showed changes in activity that track norm PE, positive PE (posPE), negative PE (negPE) or variance PE. We constructed two different design matrices: one with both norm PE and variance PE as parametric regressors (entered without orthogonalization), and one with posPE and negPE as parametric regressors (entered without orthogonalization). While the design of the task was optimized for looking for PE signals at the time of the offer reveal, during data collection, we also decided to look for these signals at the time of the response. This decision was prompted by the publication of results in rodents that showed that PE signals can also occur when the behavior/response is initiated (Syed *et al.*, 2016). Hence, we investigated the presence of PE signals related to social norms when participants saw the offers (at the offer screen) and when they were prompted to make their decision (at the response screen). Parametric regressors were convolved to these onsets in separate analyses. Regressors of interest were estimated at the subject level before being entered into second level random effect analyses. Significance thresholds at the cluster level were calculated using AFNI's 3dClustSim (post December 2015 version part of AFNI's 16.2.06 version). For whole-slab analyses (a mask of all voxels that were covered by the axial slab from at least 20 participants), we calculated that the minimum cluster size was 31 voxels for the cluster to be significant at $P < 0.05$ with an alpha of 0.005 at the voxel level. When looking at the SN/VTA complex only, the minimum cluster size threshold was six voxels to be significant at $P < 0.05$ with an alpha of 0.005 at

the voxel level. We also performed exploratory region of interest analyses on the functionally active clusters found in the GLM analyses by looking at BOLD activity patterns and signal time courses (see Supplementary material for details).

Results

Behavioral results

We first tested for behavioral effects in our participants that would indicate a norm adaptation process. We hypothesized that as in Xiang *et al.* (2013) participants previously conditioned on high vs low offers would behave differently toward subsequent medium offers. Consistent with our hypothesis, we found that, overall, participants in the H2M group were less happy (bootstrapping $P < 0.05$) with and rejected more (bootstrapping $P < 0.05$) medium offers than the L2M group (when considering the median rejection rate and mean emotional rating for 'fair' offers between 6 and 10\$ for each participant) (Figure 3). An examination of the time course of the mean rejection rate throughout the task illustrates the norm adaptation process (Figure 3 inset). We can see that during the conditioning part of the experiment, the L2M group mean rejection rate decreased with the succession of offers from the 'greedy' distribution, whereas the H2M group mean rejection rate increased with the succession of offers from the 'generous' distribution. Furthermore, for both groups, the mean rejection rate seemed to plateau at the end of the conditioning portion of the experiment (especially for the L2M group). This is the type of result we would expect if participants' internal norm gradually got close to the mean of their respective offer distribution, stabilizing their rejection rate (the decision of rejecting being governed in part by the comparison between their internal norm and the received offer). After the offer distribution changes to a medium or 'fair' value for both groups, we can see the adaptation process starting again. These results and observations are consistent with a norm adaptation process.

Modeling results

Summary statistics of the envy, guilt and temperature (see Supplementary material for details) estimates and the negative log-likelihoods of the fit of the individual choice data to our choice model outlined in the Supplementary material are presented in Table 1.

Imaging results

Having shown that participants' internal norm adapts to the received offers, we next sought to test if the brainstem was involved in processing information about social norms by looking for regions that would track norm PE, posPE, negPE or variance PE. We found a neural signature of these values in the SN/VTA and this signature was more apparent at the time when participants were prompted to respond than at the time of the offer reveal.

Results: time of the response screen. First, consider the analyses with parametric regressors time-locked to the time when participants were prompted to respond. Results from these analyses are presented in Table 2 and Figures 4 and 5. We found no region tracking norm PE at the time of the response screen within the SN/VTA complex. However, we identified several clusters within the SN/VTA complex where the signal correlated with negPE and posPE. Although there was no spatial overlap between the clusters where BOLD signal correlated with negPE and posPE (see Figure 4), directly comparing the relation

between BOLD signal and positive vs negative PE (i.e. negPE > posPE; posPE > negPE) revealed that only the clusters found in the posPE analysis showed the expected relation: a significantly stronger relation to posPE than negPE (see Table 2; Supplementary Table S2).

We also found several clusters coding for variance PE and in particular, clusters in bilateral anterior insula, and two clusters within the SN/VTA complex (Figure 5). To further investigate the pattern of activity within the SN/VTA, we conducted exploratory ROI analyses. First, the relation between BOLD activity and norm PE seemed to display the expected 'U' shape (Preuschoff *et al.*, 2008) within these clusters (Figure 5). Second, looking at the variance PE time course we can see that the brain responses to variance PE show a low < medium < high pattern and that this pattern does not seem to be a delayed response (see Supplementary material and Supplementary Figure S1) to the offer screen although because of the task design (the response screen was always 4 s after the offer screen) their contributions cannot be independently estimated. The results from these control analyses support our findings that, when we are prompted to make a decision based (at least partially) on an internal social norm, the SN/VTA processes information related to the difference between our internal norm's variance and the variance we perceive in the environment.

Results: time of the offer reveal. Results for the analyses focusing on brain signals covarying with our regressors of interest when subject saw the offer split are presented in Table 3 and Figure 6. In brief, when looking at the time of the offer reveal, no region within the brainstem was found to track norm PE, negPE or posPE. We found a small cluster within the SN/VTA complex with signal positively correlating with variance PE. An exploratory ROI analysis revealed that BOLD activity within this cluster displayed what resembles a 'U' shape relation with norm PE values. This relation was however less clear than for the clusters found at the time of the response screen. Note that we also found a cluster within the anterior insula where BOLD activity positively correlated with the variance PE signal.

Importantly, the variance PE signal was selectively present in the SN/VTA when participants saw new offers (offer screen) and when they had to choose (response screen). Indeed, it was not found when participants were informed that they were paired to a new partner (the new partner screen) (see Supplementary material for details and Supplementary Table S1).

Discussion

Using a norm adaptation paradigm and an imaging approach tailored for the study of the midbrain, we show for the first time that the human SN/VTA complex is involved in social norm processing. Indeed, our results suggest that the dopamine system codes for positive PE, negative PE and variance PE signals during an UG. These signals are theoretically used to learn about social norms and guide our actions in the face of norm violation. By probing the brain responses at different phases of the decision making process, we were able to show that these learning signals are not just present when new information about norms is processed but are also present when decisions have to be made using this new information.

Participants that repeatedly received high offers from different Proposers rejected on average more and were less pleased with subsequent medium offers than participants that repeatedly received low offers from different Proposers. This result is consistent with a norm conditioning effect—changes in

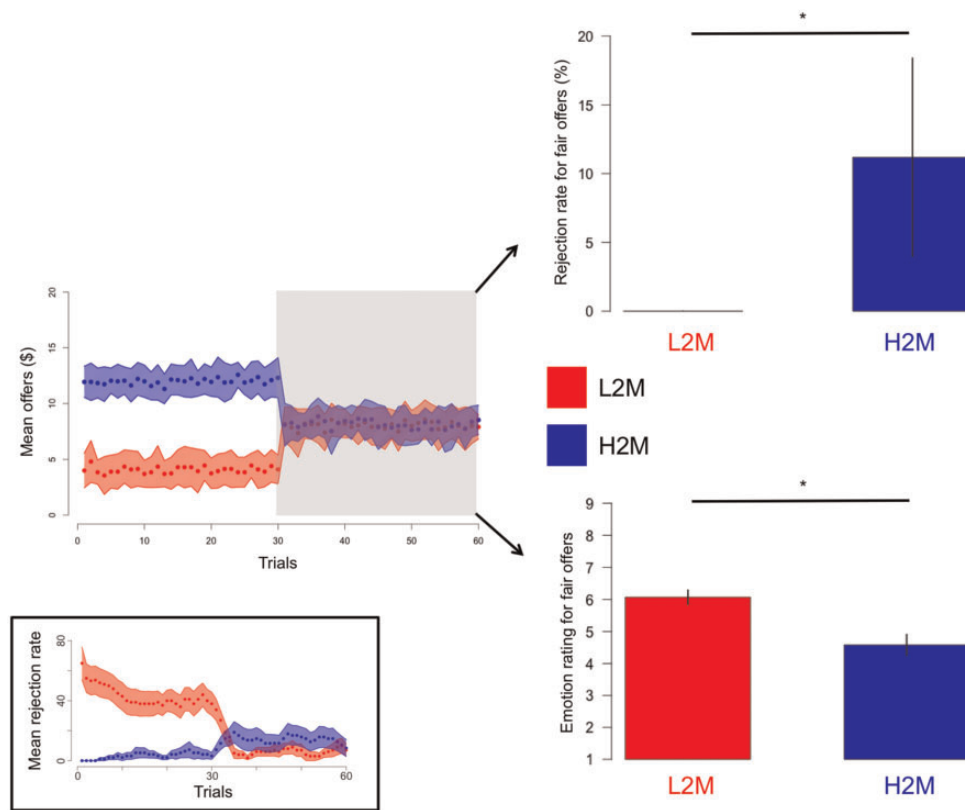


Fig. 3. Behavioral effects of the norm adjustment task. Even if during the second half of the task the L2M (red) and H2M (blue) groups both received 'fair' offers, participants from the H2M group rejected significantly more and were significantly less happy about these 'fair' offers. The grey shaded area depicts the second half of the task from which the behavioral data were extracted. The blue and red shaded area represents \pm s.d. while error bars on the bar graphs represent SEM. There are no error bars for the rejection rates of the L2M group as all the participants in this group had a point estimate (median) rejection rate for 'fair' offers between 6\$ and 10\$=0%. Inset: Mean rejection rate for the low-to-medium (red) and high-to-medium (blue) groups. Individual rejection rates were calculated using a sliding window of 10 trials. The blue and red shaded area represents \pm SEM. L2M: low-to-medium group; H2M: high-to-medium group; s.d.: standard deviation; SEM: standard error of the mean. *bootstrapping $P < 0.05$.

Table 1. Summary statistics from the behavioral modeling analysis

	Mean	s.d.
Envy	2.26	3.16
Guilt	0.34	0.44
Temperature	1.61	2.43
Negative log-likelihood	7.62	9.21

The envy and guilt parameters are used to adjust (reduce) the subjective value of offers received by the participants with higher values meaning greater sensitivity. The envy parameter (0, 10) is an estimate of the participants' sensitivity to norm violations that are disadvantageous to them (i.e. negative norm PE). The guilt parameter (0, 1) is an estimate of the participants' sensitivity to norm violations that are advantageous to them (i.e. positive norm PE). The higher the temperature parameter is (0, 10), the more diffuse and variable the participant's choices are. Negative log-likelihood is a measure of the goodness-of-fit of the model and was calculated by comparing the model prediction (using the optimized parameters for each participant) to the actual participants' choices over all trials and then averaged over participants. The negative log-likelihood value for our Bayesian observer model was significantly smaller than the one from a model where choices were made at random (bootstrapping < 0.001) and significantly smaller than a null distribution of negative log-likelihoods obtained from fitting a thousand times our Bayesian observer model to the participants' actual responses but randomly shuffled (within participants) (bootstrapping < 0.001).

normative expectations—and replicate data from a previous study which used a very similar UG task (Xiang et al., 2013). We also found very similar estimates for overall model fit, envy and guilt to the ones reported in Xiang et al. (2013). These converging

results strengthen the validity of this paradigm in studying norm adaptation and can be taken as evidence that humans can sample social information from their environment to update their internal norms and that these changes can impact their behaviors.

We were especially interested in investigating the potential role of the midbrain dopaminergic system in this norm updating and decision-making process. To do so, we looked for neural activity in the SN/VTA that covaried with norm violation signals: PE, positive PE, negative PE and variance PE. We found that the activity within the SN/VTA covaried with positive norm PE, negative norm PE and for variance PE, suggesting that the SN/VTA plays a role in norm updating and in norm based decision-making.

As mentioned before, the PE signals related to our participants' social norm is similar in theory to the classical reward PE signals that have been shown to be represented within the SN/VTA (Dreher et al., 2006; D'Ardenne et al., 2008; Pauli et al., 2015). Therefore, a parsimonious interpretation of our findings is that the SN/VTA is involved in generalized value learning and value PE coding including during social situations. However, our behavioral results argue against the possibility that participants were only computing a signal related to 'values/rewards' *per se* (e.g. the distribution of rewards in a social context). Indeed, throughout our tasks, participants decided to reject offers which resulted in them getting 0\$. They therefore chose 0\$ over a possible non-negative reward which is contrary to the behavior of a

Table 2. Summary of the results for the parametric regressors at the time of the response screen

Brain region	L/R	Voxels	Z-t	x	y	z
posPE-Whole-slab						
Medial Frontal Gyrus	L	59	3.71	-9	46	14
IFG	L	35	3.34	-50	18	1
Medial Frontal Gyrus	L	33	2.89	-1	46	25
posPE-ROI						
Midbrain-SN/VTA (VTA) ^{a,b}	L	11	3.88	-4	-26	-19
Midbrain-SN/VTA (VTA) ^b	R	10	3.75	6	-17	-9
negPE-Whole-slab						
Insula/IFG ^a	R	153	-2.93	46	20	-5
Middle Frontal Gyrus	R	50	-3.41	31	61	10
negPE-ROI						
Midbrain-SN/VTA (VTA) ^a	L	13	-3.82	-1	-23	-26
Midbrain-SN/VTA (VTA)	R	13	-2.91	4	-17	-15
variance PE-Whole-slab						
Insula/IFG ^a	R	276	2.91	42	19	-5
Middle Frontal Gyrus ^a	L	174	2.95	-47	50	6
Cerebellum-Culmen ^a	R	170	3.26	43	51	-30
ACC	L	137	3.52	-1	42	4
Cerebellum-Tuber	L	100	3.37	-43	-69	-28
Middle Frontal Gyrus	R	95	3.97	43	51	10
Insula/IFG	L	89	3.50	-29	24	-3
Midbrain-SN/VTA (VTA) ^a	L	79	3.20	-4	-17	-19
IFG ^a	R	75	3.39	43	37	-1
Fusiform Gyrus	R	37	4.05	58	-48	-17
Cerebellum-Tuber	R	36	3.70	-40	-58	-22
variance PE-ROI						
SN/VTA (SN)	L	17	3.40	-6	-13	-15
SN/VTA (SN)	L	6	3.19	-7	-21	-15

Clusters' peak voxel coordinates in brainstem normalized Talairach space and Z - t statistics for analyses done with regressors onsets locked to the response screen. t-values were transformed into Z-scores because the total number of participants that had data for specific voxels could change due to the use of individually defined acquisition slabs. Statistics were only calculated on voxels that had at least 20 participants. Whole-slab results survived FWE $P < 0.05$ corrected 31 voxels minimum cluster size (alpha of 0.005 at the voxel level). ROI results survived FWE $P < 0.05$ corrected 6 voxels minimum cluster size (alpha of 0.005 at the voxel level). L: left; R: right; ACC: anterior cingulate cortex; SN/VTA: substantia nigra/ventral tegmental area; IFG: inferior frontal gyrus; ROI: region of interest; posPE: positive norm prediction error; negPE: negative norm prediction error; variance PE: variance prediction error.

^aIndicate that these clusters were also significant at the cluster $P < 0.05$ with an alpha of 0.001 at the voxel level (minimum cluster size: 14 voxels for whole-slab; 3 voxels for ROI).

^bIndicate that the cluster overlaps with some clusters found by the contrast posPE > negPE (see Supplementary material).

purely reward-seeking agent and more in line with a costly punishment behavior related to social norm enforcement (Fehr and Gächter, 2002; Henrich et al., 2006; Boyd et al., 2010). The behavior observed in our UG—rejection of non-negative rewards and even of offers that were higher than 10\$ (hyperfair)—also suggests that our participants really thought they were in a social context. Furthermore, the conditioning effect we observed where the rejection pattern of participants adapted to the changes in the offer distribution cannot be explained by a purely reward learning process: changes in rejection rate are better explained by changes in the normative expectations of participants—what they think the Proposer *ought* to propose (Bicchieri, 2016). In other words, their behavior suggests that they considered the experienced offer distributions as normative—not as a purely descriptive reward distribution—and therefore updated their own internal social norm

accordingly. This being said, since we did not use a non-social control condition, our paradigm does not allow us to study if and how our norm PE signals are different/similar from/to reward PEs or to formally test if these signals would be present if the participants played with a non-human partner. Nevertheless, our novel results showing that the SN/VTA is involved in detecting social norm deviations and in updating social norms importantly extend the role of the midbrain dopamine system from 'simple' reward-based decision-making to social decision-making.

Data from animal (Fiorillo et al., 2003) and human (Dreher et al., 2006; Preuschoff et al., 2006) studies have shown that activity in the SN/VTA varies as a function of uncertainty/risk but no study has specifically looked at its potential role in the updating/learning of this estimate. We present evidence that the SN/VTA is implicated in the computing of error signals used to update our estimate of the uncertainty about our internal norm. Indeed, clusters within the SN/VTA covaried with the trial-by-trial value of the variance PE signal. In addition, we found that the BOLD activity within these clusters displayed the expected 'U' shape in relation to the norm PE values. Coding for risk/variance PE has been shown previously to be supported by the insula (Preuschoff et al., 2008; d'Acremont et al., 2009; Xiang et al., 2013), a result confirmed by our own data.

Interestingly, we found that the SN/VTA and insula coded for the variance PE at the time when participants were prompted to choose between accepting and rejecting the offers (response screen) in addition to the offer reveal time. A variance PE signal at the time of the offer split was expected as it is the first moment when expectations about norms' variance can be compared to new information. The origin and function of variance PE signal at the time of the response screen is less clear since previous studies on uncertainty PE have not looked at this time point (Preuschoff et al., 2008; d'Acremont et al., 2009; Xiang et al., 2013). Nevertheless, it is unlikely that these signals are only an artifact of our task structure. First, we showed that in addition to varying linearly with the value of the variance PE, the signal within two out of three SN/VTA clusters identified at the time of the response screen displayed the hypothesized 'U' shape (Preuschoff et al., 2008) with the norm PE values. Second, previous work has shown that the risk PE signal (within the insula in this case) is fast and brief (Preuschoff et al., 2008) suggesting that differences in BOLD activity between high, medium and low signals should be apparent soon after the onset of the event. Looking at the BOLD time course analysis (from which the motor behavior time series was regressed out), we can see that the signal starts to clearly separate at the time of the response screen, not the time of the offer split. This suggests that the deflection we observed in the BOLD signal at the time of the response is not a delayed response to the offer split revelation. Third, while the variance PE responses was present at the offer and response screen times, it was absent at the start of the next trial. These results suggest that the BOLD signal we measured is changing in relation to the variance PE value on a trial-to-trial basis and probably plays a role in norm learning and decision-making. To get a clearer picture of the role of the variance PE at the time of the decision, future work will need to systematically assess how this signal is created and used by the brain and investigate the role of the insula and the SN/VTA at different times in the decision-making process.

Our results suggest that the SN/VTA does not process norm PE as a continuum but rather separately tracks positive and negative PE. Indeed, we found no region in the midbrain with activity linearly varying with norm PE when modeled as a

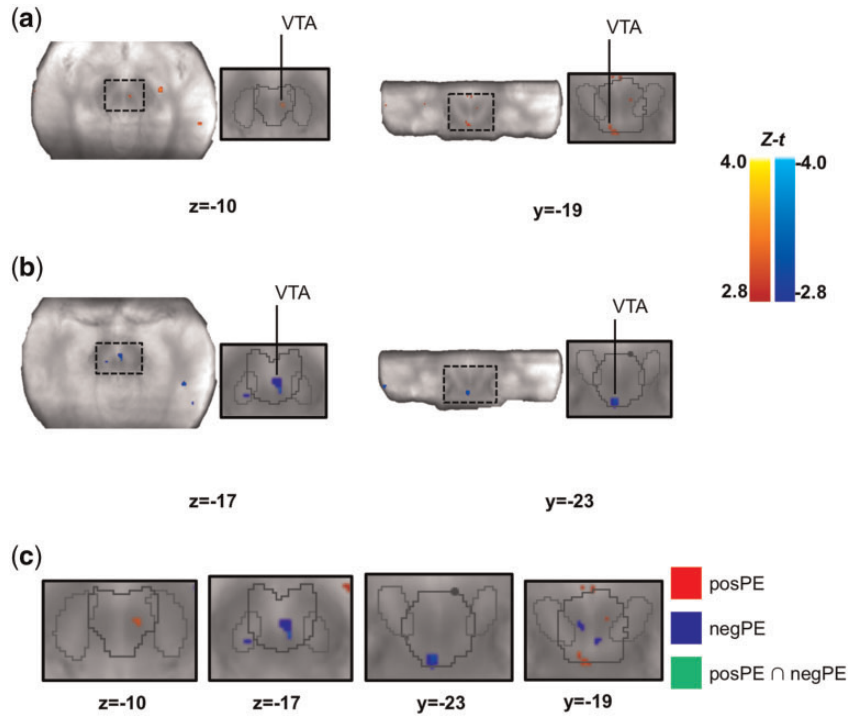


Fig. 4. Brain activity correlating with the norm PE signal at the time of the response. (A, B) We separately looked at the brain activity correlating with the positive (A) and negative (B) components of the norm PE. The identified SN/VTA clusters are displayed on the mean proton-density image. (C) We compared the maps from the positive and negative norm PEs analyses using the approach presented in Nichols et al. (2005) and found no overlap. Note that no voxel in the figure shows a conjunction (green color). The SN/VTA cluster is displayed on the mean proton-density image. The SN (light grey) and VTA (dark grey) are outlined based on the Murty et al. (2014) atlas. Images were thresholded at $P < 0.005$ uncorrected for display purposes. t -values were transformed into Z-scores because the total number of participants that had data for specific voxels could change due to the use of individually defined acquisition slabs. Statistics were only calculated on voxels that had at least 20 participants. SN/VTA: substantia nigra/ventral tegmental area.

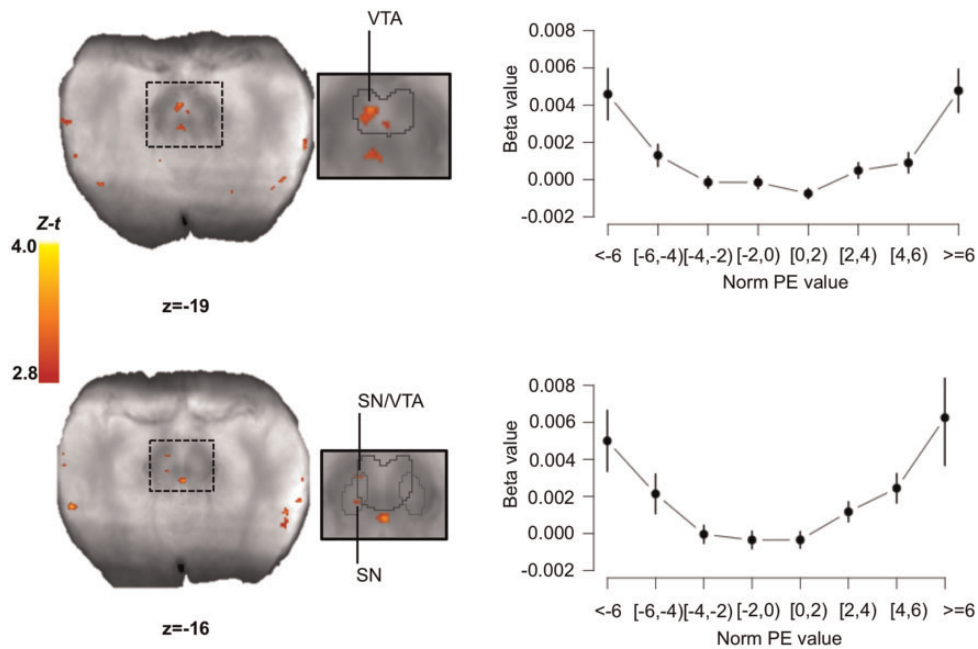


Fig. 5. Brain activity correlating with the variance PE signal at the time of the response. We found three clusters within the SN/VTA which BOLD signal correlated with the value of the variance PE. Using these clusters as ROIs, we extracted the mean beta values for different binned-values of norm PEs (in \$). Right plots show the 'U' shape relationship between brain activity expressed as beta values and the values of the norm PE for top: cluster (-4 -17 -19); for bottom most rostral cluster (-6 -13 -15). For the bottom most caudal cluster at (-7 -21 -15), the relation did not display the 'U' shape relationship. Error bars represent SEM. The SN/VTA clusters are displayed on the mean proton-density image. The SN (light grey) and VTA (dark grey) are outlined based on the Murty et al. (2014) atlas. Images were thresholded at $P < 0.005$ uncorrected for display purposes. t -values were transformed into Z-scores because the total number of participants that had data for specific voxels could change due to the use of individually defined acquisition slabs. Statistics were only calculated on voxels that had at least 20 participants. SN/VTA: substantia nigra/ventral tegmental area; norm PE: norm prediction error; SEM: standard error of the mean.

continuum. However, when considering positive and negative norm PE as separate signals we identified spatially distinct regions of the SN/VTA coding for each signal. Furthermore, the fact that we found a negative correlation between the BOLD signal and the negative norm PE (i.e. increased activity with increasing absolute negative values) suggests that the SN/VTA treats similarly positive and negative norm PE: the larger the deviation from the norm, the greater the increase in BOLD activity. Xiang et al. (2013) who used a very similar task, also found separate processing for positive and negative norm PE in the striatum which has important bidirectional connections with SN/VTA (Yetnikoff et al., 2014), suggesting that signaling for norm deviation is similar in various parts of the dopamine system. Interestingly, Xiang et al. (2013) responses were at the time of the offer split (they only tested this time point) whereas we only found positive and negative norm PE signal in the SN/VTA at the response screen. The striatum is an important projection target for midbrain dopamine neurons (Björklund and Dunnett, 2007) and one could have expected that the activity seen in the striatum at the time of the offer split would originate from the

Table 3. Summary of the results for the parametric regressors at the time of the offer screen

Brain region	L/R	Voxels	Z-t	x	y	z
negPE-Whole-slab						
Insula/IFG	R	24	-3.21	34	18	-7
variance PE-Whole-slab						
Insula ^a	R	93	3.07	34	18	-7
variance PE-ROI						
SN/VTA (VTA)	L	10	3.29	-2	-13	-22

Clusters' peak voxel coordinates in brainstem normalized Talairach space and Z-t statistics for analyses done with regressors onsets locked to the offer screen. t-values were transformed into Z-scores because the total number of participants that had data for specific voxels could change due to the use of individually defined acquisition slabs. Statistics were only calculated on voxels that had at least 20 participants. Whole-slab results survived FWE $P < 0.05$ corrected 31 voxels minimum cluster size (alpha of 0.005 at the voxel level). ROI results survived FWE $P < 0.05$ corrected six voxels minimum cluster size (alpha of 0.005 at the voxel level). L: left; R: right; IFG: inferior frontal gyrus; SN/VTA: substantia nigra/ventral tegmental area; ROI: region of interest; negPE: negative norm prediction error; variance PE: variance prediction error.

^aIndicate that these clusters were also significant at the cluster $P < 0.05$ with an alpha of 0.001 at the voxel level (minimum cluster size: 14 voxels for whole-slab; 3 voxels for ROI).




Fig. 6. Brain activity correlating with the variance PE signal at the time of the offer. We found one cluster within the SN/VTA which BOLD signal correlated with the value of the variance PE. Using this cluster as a ROI we extracted the mean beta values for different binned-values of norm PEs (in \$). Right plots show the relationship between brain activity expressed as beta values and the values of the norm PE. Error bars represent SEM. The SN/VTA cluster is displayed on the mean proton-density image. The SN (light grey) and VTA (dark grey) are outlined based on the Murty et al. (2014) atlas. Images were thresholded at $P < 0.005$ uncorrected for display purposes. t-values were transformed into Z-scores because the total number of participants that had data for specific voxels could change due to the use of individually defined acquisition slabs. Statistics were only calculated on voxels that had at least 20 participants. SN/VTA: substantia nigra/ventral tegmental area; norm PE: norm PE SEM: standard error of the mean.

SN/VTA. Alternatively, it is possible that the information about norm PE received by the ventral striatum originates from regions such as the mOFC or the anterior insula that were also shown to code for positive and negative norm PE (Xiang et al., 2013). In turn, since the striatum also sends projections to the midbrain dopamine neurons (Yetnikoff et al., 2014), it is possible that the activity we found in the SN/VTA at the time of the response screen comes from information conveyed by the ventral striatum. But since our imaging procedure precluded us from acquiring data from the striatum, we could not check for possible activation in the striatum at the time of the response screen.

In several models of choice behavior in the UG (Fehr and Schmidt, 1999; Bicchieri, 2005), the computation of the values assigned to each option (accept/reject) includes the difference between the expected value and the actual offer, and thus takes into account a form of PE. Therefore, the negative and positive PE signals we observed in the SN/VTA at the time of the response screen may be indicative of the computation of choice values used during the decision process. Interestingly, a recent study in rodents has suggested that when actions have to be inhibited (delayed) before the reward is received such as in a No-Go condition, reward PE can be observed when the animals initiate the action to get the reward, well after the reward cue has been seen (when we would expect the reward PE signal to arise) (Syed et al., 2016). This could explain why we observed the norm PE error (and the variance PE) signals at the response time as our participants had to initiate their actions (i.e. response) 4 s after the cue (i.e. offer reveal). Importantly, these hypotheses are tentative and future studies will be needed to uncover the functional role of these signals at the time when a response to a decision has to be made.

Our finding of spatially distinct regions coding for positive and negative norm PE within the SN/VTA is of significance for several reasons. First, if there is strong evidence that the SN/VTA encodes positive/appetitive signals (Dreher et al., 2006; D'Ardenne et al., 2008), its role in processing aversive stimuli remain a source of debate (Fiorillo, 2013; Proulx et al., 2014) with data mainly coming from rodents or non-human primates studies. Recent human studies have however offered evidence that the SN and VTA are active during the processing of cues predicting aversive or negative sensory stimuli (electric shocks; Hennigan et al., 2015) and that appetitive as well as aversive PE in the gustatory domain are represented in distinct regions of the SN/VTA complex (Pauli et al., 2015). Our results thus add to

the view that this structure is involved in the processing of both positive and negative stimuli and extend this role into the social norm domain. It is important to point out that the only clusters that showed a significantly stronger relation between BOLD activity and posPE vs negPE were the ones that were tracking posPE. This may suggest that the functional specificity of the cluster we identified as tracking negPE (i.e. ONLY coding for negPE) is lower than the one for the clusters coding for posPE. Second, and more generally, while the SN/VTA's role in learning in the monetary and gustatory domains is well known (Dreher et al., 2006; D'Ardenne et al., 2008, 2013; Pauli et al., 2015), by finding that it tracks positive and negative (and variance) PE in a task where participants update their norm, we show that the SN/VTA is also involved in social-norms-based learning and decision-making.

Conclusions

Our results show that the SN/VTA encodes PE signals related to social norms and their uncertainty when individuals receive new information about their social environment and when they are prompted to make a decision based on this new information. This suggests that midbrain dopamine neurons are not only part of the network that enables us to track and update social norms, detect deviations from these norms, but are also involved with how we respond to such deviations.

Supplementary data

Supplementary data are available at SCAN online.

Acknowledgements

Sebastien Hetu is now at Université de Montréal. Yi Luo is now at Beijing Normal University. Kimberlee D'Ardenne is now at Arizona State University. We thank Alec Solway and Iris Vilares for their helpful comments on previous versions of the manuscript.

Funding

This study was supported by a National Institute on Drug Abuse (NIDA) grant R01 DA11723, a National Institute of Mental Health (NIMH) grant 5R01MH085496-09. P.R.M. is supported by a Principal Research Fellowship from The Wellcome Trust. S.H. is supported by Fellowships from the Fonds de Recherche Santé (FRQS) and Canadian Institute of Health Research (CIHR).

Conflict of interest. None declared.

References

Bach, D.R., Dolan, R.J. (2012). Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature Reviews Neuroscience* 13(9), 572–86.

Beissner, F. (2015). Functional MRI of the brainstem: common problems and their solutions. *Clinical Neuroradiology* 25(2), 251–7.

Bicchieri, C. (2005) *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge, MA: Cambridge University Press.

Bicchieri, C. (2016) *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford: Oxford University Press.

Björklund, A., Dunnett, S.B. (2007). Dopamine neuron systems in the brain: an update. *Trends in Neurosciences* 30(5), 194–202.

Boyd, R., Gintis, H., Bowles, S. (2010). Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science* 328(5978), 617–20.

Camerer, C. (2003) *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.

Chang, L.J., Sanfey, A.G. (2013). Great expectations: neural computations underlying the use of social norms in decision-making. *Social Cognitive and Affective Neuroscience* 8(3), 277–84.

Cox, R.W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research* 29(3), 162–73.

d'Acromont, M., Lu, Z.-L., Li, X., Van der Linden, M., Bechara, A. (2009). Neural correlates of risk prediction error during reinforcement learning in humans. *NeuroImage* 47(4), 1929–39.

D'Ardenne, K., McClure, S.M., Nystrom, L.E., Cohen, J.D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science* 319(5867), 1264–7.

D'Ardenne, K., Lohrenz, T., Bartley, K., Montague, P.R. (2013). Computational heterogeneity in the human mesencephalic dopamine system. *Cognitive, Affective, and Behavioral Neuroscience* 13(4), 747–56.

Dagli, M.S., Ingeholm, J.E., Haxby, J.V. (1999). Localization of cardiac-induced signal change in fMRI. *NeuroImage* 9(4), 407–15.

Dreher, J.-C., Kohn, P., Berman, K.F. (2006). Neural coding of distinct statistical properties of reward information in humans. *Cerebral Cortex* 16(4), 561–73.

Eapen, M., Zald, D.H., Gatenby, J.C., Ding, Z., Gore, J.C. (2011). Using high-resolution MR imaging at 7T to evaluate the anatomy of the midbrain dopaminergic system. *American Journal of Neuroradiology* 32(4), 688–94.

Enzmann, D.R., Pelc, N.J. (1992). Brain motion: measurement with phase-contrast MR imaging. *Radiology* 185(3), 653–60.

Fehr, E., Gächter, S. (2002). Altruistic punishment in humans. *Nature* 415(6868), 137–40.

Fehr, E., Schmidt, K.M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics* 114(3), 817–68.

Fiorillo, C.D. (2013). Two dimensions of value: dopamine neurons represent reward but not aversiveness. *Science* 341(6145), 546–9.

Fiorillo, C.D., Tobler, P.N., Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299(5614), 1898–902.

Gu, X., Wang, X., Hula, A., et al. (2015). Necessary, yet dissociable contributions of the insular and ventromedial prefrontal cortices to norm adaptation: computational and lesion evidence in humans. *The Journal of Neuroscience* 35(2), 467–73.

Hennig-Schmidt, H., Li, Z.-Y., Yang, C. (2008). Why people reject advantageous offers—non-monotonic strategies in ultimatum bargaining: evaluating a video experiment run in PR China. *Journal of Economic Behavior & Organization* 65(2), 373–84.

Hechter, M., Opp, K.D. (2001) *Social norms*. New York: Russell Sage Foundation.

Hennigan, K., D'Ardenne, K., McClure, S.M. (2015). Distinct mid-brain and habenula pathways are involved in processing aversive events in humans. *The Journal of Neuroscience* 35(1), 198–208.

Henrich, J., McElreath, R., Barr, A., et al. (2006). Costly punishment across human societies. *Science* 312(5781), 1767–70.

Klucharev, V., Hytönen, K., Rijpkema, M., Smidts, A., Fernández, G. (2009). Reinforcement learning signal predicts social conformity. *Neuron* 61(1), 140–51.

Lang, P. (1980) Behavioral treatment and bio-behavioral assessment: computer applications. In: Sidowski, J.B., Johnson, J.H.,

- Williams, T.A., editors. *Technology in Mental Health Care Delivery Systems*. Norwood, NJ: Ablex.
- Lawson, R.P., Drevets, W.C., Roiser, J.P. (2013). Defining the habenula in human neuroimaging studies. *NeuroImage* **64**, 722–7.
- Ma, W.J., Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience* **37**(1), 205–20.
- McClure, S.M., Berns, G.S., Montague, P.R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron* **38**(2), 339–46.
- Montague, P.R., Dayan, P., Sejnowski, T.J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience* **16**(5), 1936–47.
- Montague, P.R., Lohrenz, T. (2007). To detect and correct: norm violations and their enforcement. *Neuron* **56**(1), 14–8.
- Mooney, C.Z., Duval, R.D., Duval, R. (1993) *Bootstrapping: A Nonparametric Approach to Statistical Inference*. Sage.
- Murty, V.P., Shermohammed, M., Smith, D.V., Carter, R.M., Huettel, S.A., Adcock, R.A. (2014). Resting state networks distinguish human ventral tegmental area from substantia nigra. *NeuroImage* **100**, 580–9.
- Nichols, T., Brett, M., Andersson, J., Wager, T., Poline, J.-B. (2005). Valid conjunction inference with the minimum statistic. *NeuroImage* **25**(3), 653–60.
- 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**(2), 329–37.
- Pauli, W.M., Larsen, T., Collette, S., Tyszka, J.M., Seymour, B., O’Doherty, J.P. (2015). Distinct contributions of ventromedial and dorsolateral subregions of the human substantia nigra to appetitive and aversive learning. *The Journal of Neuroscience* **35**(42), 14220–33.
- Preusschoff, K., Bossaerts, P., Quartz, S.R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* **51**(3), 381–90.
- Preusschoff, K., Quartz, S.R., Bossaerts, P. (2008). Human insula activation reflects risk prediction errors as well as risk. *The Journal of Neuroscience* **28**(11), 2745–52.
- Proulx, C.D., Hikosaka, O., Malinow, R. (2014). Reward processing by the lateral habenula in normal and depressive behaviors. *Nature Neuroscience* **17**(9), 1146–52.
- Rilling, J.K., Sanfey, A.G., Aronson, J.A., Nystrom, L.E., Cohen, J.D. (2004). The neural correlates of theory of mind within interpersonal interactions. *NeuroImage* **22**(4), 1694–703.
- Rolls, E.T., McCabe, C., Redoute, J. (2008). Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task. *Cerebral Cortex* **18**(3), 652–63.
- Sanfey, A.G., Stallen, M., Chang, L.J. (2014). Norms and expectations in social decision-making. *Trends in Cognitive Sciences* **18**(4), 172–4.
- Schultz, W., Dayan, P., Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* **275**(5306), 1593–9.
- Starmans, C., Sheskin, M., Bloom, P. (2017). Why people prefer unequal societies. *Nature Human Behaviour* **1**, 0082.
- Syed, E.C., Grima, L.L., Magill, P.J., Bogacz, R., Brown, P., Walton, M.E. (2016). Action initiation shapes mesolimbic dopamine encoding of future rewards. *Nature Neuroscience* **19**(1), 34–6.
- Van’t Wout, M., Kahn, R.S., Sanfey, A.G., Aleman, A. (2006). Affective state and decision-making in the ultimatum game. *Experimental Brain Research* **169**(4), 564–8.
- Vilares, I., Kording, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences* **1224**(1), 22–39.
- Xiang, T., Lohrenz, T., Montague, P.R. (2013). Computational substrates of norms and their violations during social exchange. *The Journal of Neuroscience* **33**(3), 1099–108.
- Yetnikoff, L., Lavezzi, H.N., Reichard, R.A., Zahm, D.S. (2014). An update on the connections of the ventral mesencephalic dopaminergic complex. *Neuroscience* **282**, 23–48.