NeuroImage 174 (2018) 35–43

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Oxytocin attenuates trust as a subset of more general reinforcement learning, with altered reward circuit functional connectivity in males

Jaime S. Ide [a], Sanja Nedic [a], Kin F. Wong [a], Shmuel L. Strey [a], Elizabeth A. Lawson [b,f], Bradford C. Dickerson [c,e,f], Lawrence L. Wald [c,f], Giancarlo La Camera [d], Lilianne R. Mujica-Parodi [a,c,f,*]

[a] Department of Biomedical Engineering, Stony Brook University School of Medicine, Stony Brook, NY, 11794, USA
[b] Neuroendocrine Unit, Massachusetts General Hospital, Boston, MA, 02114, USA
[c] Department of Radiology, A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, 02129, USA
[d] Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, NY, 11794, USA
[e] Department of Neurology, Massachusetts General Hospital, Boston, MA, 02114, USA
[f] Harvard Medical School, Boston, MA, 02115, USA

ARTICLE INFO

Keywords:
Oxytocin
fMRI
Bayesian modeling
Reinforcement learning
Human
Male

ABSTRACT

Oxytocin (OT) is an endogenous neuropeptide that, while originally thought to promote trust, has more recently been found to be context-dependent. Here we extend experimental paradigms previously restricted to *de novo* decision-to-trust, to a more realistic environment in which social relationships evolve in response to iterative feedback over twenty interactions. In a randomized, double blind, placebo-controlled within-subject/crossover experiment of human adult males, we investigated the effects of a single dose of intranasal OT (40 IU) on Bayesian expectation updating and reinforcement learning within a social context, with associated brain circuit dynamics. Subjects participated in a neuroeconomic task (*Iterative Trust Game*) designed to probe iterative social learning while their brains were scanned using ultra-high field (7T) fMRI. We modeled each subject's behavior using Bayesian updating of belief-states ("willingness to trust") as well as canonical measures of reinforcement learning (*learning rate*, *inverse temperature*). Behavioral trajectories were then used as regressors within fMRI activation and connectivity analyses to identify corresponding brain network functionality affected by OT. Behaviorally, OT reduced feedback learning, without bias with respect to positive versus negative reward. Neurobiologically, reduced learning under OT was associated with muted communication between three key nodes within the reward circuit: the *orbitofrontal cortex, amygdala,* and *lateral (limbic) habenula*. Our data suggest that OT, rather than inspiring feelings of generosity, instead attenuates the brain's encoding of prediction error and therefore its ability to modulate pre-existing beliefs. This effect may underlie OT's putative role in promoting what has typically been reported as 'unjustified trust' in the face of information that suggests likely betrayal, while also resolving apparent contradictions with regard to OT's context-dependent behavioral effects.

Introduction

Oxytocin (OT) is an endogenous neuropeptide that, when exogenously administered intranasally, has been reported to increase people's willingness to trust other humans (Kosfeld et al., 2005), even after betrayal (Baumgartner et al., 2008). The dominant hypothesis is that OT increases trust by reducing fear and associated brain activations in the *amygdala*, *midbrain*, and *dorsal striatum* (Baumgartner et al., 2008). Supporting this hypothesis are findings that OT attenuates the response of the amygdala, as well as that of its related circuits, to fear (Kirsch et al., 2005), conditioned fear (Petrovic et al., 2008), and fearful faces (Domes et al., 2007; Gamer et al., 2010).

The first two studies reporting the behavioral (Kosfeld et al., 2005) and neural (Baumgartner et al., 2008) effects of oxytocin in humans used versions of a neuroeconomic task known as the *Trust Game,* in which player *A* makes a decision about how to split money with player *B*, and then *B* does the same with *A*. Thus, *A*'s split reflects assumptions ('trust') about *B*'s predicted reciprocal behavior. OT increases generosity in the

* Corresponding author. Department of Biomedical Engineering, Stony Brook University, Stony Brook, NY, 11794-5281, USA.
E-mail address: lilianne.strey@stonybrook.edu (L.R. Mujica-Parodi).

https://doi.org/10.1016/j.neuroimage.2018.02.035
Received 11 May 2017; Received in revised form 24 January 2018; Accepted 17 February 2018
Available online 25 February 2018

Trust Game but not in the simpler Dictator Game, the latter of which eliminates assumptions of reciprocity (Zak et al., 2007). Specifically, subjects who were given OT did not change their trusting behavior after receiving information that many trustees had betrayed their trust in previous interactions, whereas subjects who received placebo reduced their trusting behavior after being so informed (Baumgartner et al., 2008). Importantly, in the initial (Kosfeld et al., 2005) study, the effect was reported to be highly trust-specific: oxytocin did not change the behavior of trustees in the Trust Game, nor the behavior of investors in a risky decision task not involving trust. While neuroimaging reports on OT have focused almost exclusively upon the neuropeptide's effect on limbic regions typically associated with fear, one of the earliest fMRI papers on OT showed that it also reduces activity in the *bilateral caudate* (Baumgartner et al., 2008), a key brain structure in reward learning. This raises the question of whether OT's putative effect in blocking the effects of aversive or aversively conditioned stimuli might actually be consequent to a more general diminished recruitment of the reward learning circuit, with associated diminished behavioral adaptation to information feedback.

Previous studies have extensively examined effects of OT in terms of the initial instinct to trust or fear in the absence of (known) prior information (Kirsch et al., 2005; Kosfeld et al., 2005; Baumgartner et al., 2008; Petrovic et al., 2008). However, social relationships typically evolve over time, in response to iterative feedback over the course of many interactions. Therefore, in order to probe the brain circuit dynamics underlying an individual's interaction-evolution, we had subjects play a multi-round version (King-Casas et al., 2005) of the Trust Game (Camerer, 2003) while undergoing 7T fMRI optimized for time-series dynamics at the single-subject level (DeDora et al., 2016). We then modeled their behavior using two approaches.

First, Bayesian modeling described dynamically evolving expectations with regard to positive outcomes. These expectation dynamics were then used as regressors for brain data, to identify neural regions of interest (Yu and Cohen, 2009; Ide et al., 2013) associated with 'trust'. A subset of these neural regions comprised a reduced functional circuit: *amygdala*, *nucleus accumbens*, *orbitofrontal cortex*, previously established by the animal (Dayan and Balleine, 2002), human (O'Doherty et al., 2003), and computational neuroscience (Dayan and Abbott, 2005) literature to underlie reinforcement learning.

Second, we assessed the degree to which subjects (Investors) learned in response to their presumed partners' (Trustees') behavior. This was done using both a simple intuitive measure of previous-trial reciprocity ('tit-for-tat'), as well as a more rigorous reinforcement learning model quantifying *exploration* (often described as 'inverse temperature,' a measure of risk-taking) and *exploitation* (the tendency to capitalize on detected patterns/rules) (Dayan and Abbott, 2005). Using psychophysiological interaction analyses (Gitelman et al., 2003) we then identified condition-specific brain connectivity within the reinforcement learning circuit.

Methods and materials

Subjects and screening procedures

Seventeen healthy male subjects ($\mu_{age} = 25.4 \pm 3.7$ years, $\mu_{weight} = 74 \pm 10$ kg, 2 left-handed) participated in a randomized double-blind within-subject/crossover experiment using a single dose intranasal oxytocin (40 IU) compared to placebo. After an initial phone screening, a study physician obtained written consent from each subject, who then underwent a History and Physical exam. Exclusion criteria included neurological/psychiatric diagnoses, body mass index >30, blood pressure >140/90 mm Hg (or controlled with medication), smoking, and nasal obstruction. Subjects were instructed to abstain from caffeine and alcohol on the day of the scan. Protocols described here were approved by the Institutional Review Boards of Stony Brook University and Partners HealthCare; all subjects provided informed consent.

Administration of oxytocin and placebo

Syntocinon (Oxytocin) Nasal Spray® (Novartis) was administered under FDA IND # 112931. Subjects received 10 sprays (40IU, 1 mL) 60 min prior to the fMRI. Placebo, identical in preparation except for the oxytocin component, was administered in the same manner in a double blind, single-dose, randomized procedure counterbalanced for order. To avoid bleed-through between conditions while controlling for order effects, each session was either oxytocin (OT) only or placebo (PL) only, conducted on separate days; OT and PL were administered at the same time on both days to control for possible diurnal variations in endogenous OT. The number of days between the two sessions ranged between 1 (for 4 out of 17 subjects) and 71, with the median being 7 ($\mu = 14$, s.d. = 20.8).

Studies looking at the effects of intranasal administration of Oxytocin have primarily used a single dose between 24 and 40 IU (Kendrick et al., 2016), with reported dosages ranging from 2 to 40 IU (Wigton et al., 2015), and dose-dependent effects observed in several studies (Cardoso et al., 2013; Quintana et al., 2017), even for lower dosages (Quintana et al., 2015, 2016). The only study to establish that intranasally-delivered neuropeptides do, in fact, cross the blood-brain barrier (Born et al., 2002), used larger doses of a closely related neuropeptide, vasopressin, at 40 and 80IU. They found that CSF concentrations began to rise within 10 min of intranasal administration and continued to increase for up to 80 min after administration. Based upon these results, we chose both the dosage and timing of the study design.

Magnetic resonance imaging

All MRI data were acquired on a 7T Siemens Magnetom scanner (32-channel head-coil array) at the Martinos Center for Biomedical Imaging at MGH. We obtained whole brain EPI BOLD data using parameters previously optimized on this scanner for dynamic fidelity of single-subject time-series (DeDora et al., 2016): SMS slice acceleration factor = 5, GRAPPA acceleration = 2, TR = 802 ms, TE = 20 ms, flip angle = 33°, 2 × 2 × 1.5 mm voxels, 748 measurements (~10 min). Field map images were acquired using: TR = 723 ms, TE1/TE2 = 4.60/5.62 ms, flip angle = 36°, and 1.7 × 1.7 × 1.5 mm voxels. T1-weighted structural volumes were acquired using a conventional MEMPRAGE sequence with 1 mm isotropic voxels and four echoes with TE1/TE2/TE3/TE4 = 1.61/3.47/5.33/7.19 ms, TR = 2530 ms, flip angle = 7°, GRAPPA acceleration = 2.

Iterative trust game

We adapted an iterative version of the Trust Game (King-Casas et al., 2005). During each scan (OT and PL), the subject ("Investor") played 20 rounds with the same opponent ("Trustee"). At the start of each round, the subject was given 20 monetary units (MU) and told to invest any amount between 0 and 20 with the Trustee. This invested amount was then tripled. The Trustee then repaid some portion of the total (0–60 MU) back to the Investor. While, in reality, the 'Trustee' was a computer-generated algorithm, subjects were told they were playing with a human; the deception was revealed following completion of the study.[1] To ensure that the 'Trustee' algorithm mimicked human behavior, parameters were estimated from data (N = 48) obtained from a previous study (King-Casas et al., 2005); randomness was set at 10%. As illustrated in Fig. 1a–b, each round consisted of a *cue to invest* (I1), *investment period* (I2), *delay*, *investment reveal* (I3), *delay*, *cue to repay* (R1), *repayment period* (R2), *delay*, *repayment reveal* (R3), *delay*, *totals reveal*, and *inter-round delay*. Delay periods were jittered between 2 and 7s using

[1] During our debriefing, prior to revealing the deception, we asked subjects about their perception of the game. None of the subjects showed evidence of questioning the cover story.