

Functional neurochemical imaging of the human striatal cholinergic system during reversal learning

Article

Accepted Version

Bell, T. ORCID: <https://orcid.org/0000-0002-9591-706X>,
Lindner, M., Mullins, P. G. and Christakou, A. ORCID:
<https://orcid.org/0000-0002-4267-3436> (2018) Functional
neurochemical imaging of the human striatal cholinergic
system during reversal learning. *The European Journal of
Neuroscience*, 47 (10). pp. 1184-1193. ISSN 0953-816X doi:
10.1111/ejn.13803 Available at
<https://centaur.reading.ac.uk/74591/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1111/ejn.13803>

Publisher: Blackwell Publishing

All outputs in CentAUR are protected by Intellectual Property Rights law,
including copyright law. Copyright and IPR is retained by the creators or other
copyright holders. Terms and conditions for use of this material are defined in
the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

DR. ANASTASIA CHRISTAKOU (Orcid ID : 0000-0002-4267-3436)

Article type : Special Issue Article

Reviewers: Steve Williams, University of Manchester, UK
John O'Doherty, California Institute of Technology, USA

Special issue on "Cholinergic Mechanisms in Adaptive Behaviour"

Functional neurochemical imaging of the human striatal cholinergic system during reversal learning

Tiffany Bell¹, Michael Lindner¹, Paul Gerald Mullins², Anastasia Christakou¹

1. School of Psychology and Clinical Language Sciences, and Centre for Integrative Neuroscience and Neurodynamics, University of Reading
2. School of Psychology, Bangor University

Correspondence to: Dr Anastasia Christakou
School of Psychology and Clinical Language Sciences
University of Reading
Whiteknights
Reading RG6 6AL
United Kingdom
Email: anastasia.christakou@reading.ac.uk

Running title: Cholinergic fMRS in reversal learning

Keywords: magnetic resonance spectroscopy; choline; striatum; reversal learning

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/ejn.13803

This article is protected by copyright. All rights reserved.

ABSTRACT

Animal studies have shown that acetylcholine (ACh) levels in the dorsal striatum play a role in reversal learning. However, this has not been studied in humans due to a lack of appropriate non-invasive techniques. Proton magnetic resonance spectroscopy (^1H -MRS) can be used to measure metabolite levels in humans *in vivo*. Although it cannot be used to study ACh directly, ^1H -MRS can be used to study choline, an ACh precursor which is linked to activity-dependent ACh release. The aim of this study was to use functional- ^1H -MRS (fMRS) to measure changes in choline levels in the human dorsal striatum during performance of a probabilistic reversal learning task. We demonstrate a task-dependent decrease in choline, specifically during reversal, but not initial, learning. We interpret this to reflect a sustained increase in ACh levels, which is in line with findings from the animal literature. This task-dependent change was specific to choline and was not observed in control metabolites. These findings provide support for the use of fMRS in the *in vivo* study of the human cholinergic system.

INTRODUCTION

Studies in rodents have demonstrated that cholinergic activity in the associative dorsal striatum (DS) is crucial for reversal learning. For example, disruption of cholinergic signalling in this region, due mainly to striatal cholinergic interneuron (CIN) activity, has been shown to impair reversal learning, whilst leaving initial learning intact (Ragozzino *et al.*, 2002, 2009; Tzavos *et al.*, 2004; McCool *et al.*, 2008; Brown *et al.*, 2010; Bradfield *et al.*, 2013). This impairment is specifically related to regressive errors after the reversal has been identified (rather than to identifying the occurrence of the reversal *per se*), suggesting interference between old and new learning (Bradfield *et al.*, 2013). Additionally, acetylcholine (ACh) efflux has been shown to increase during reversal learning, but not during initial learning (Ragozzino & Choi, 2004; Ragozzino *et al.*, 2009; Brown *et al.*, 2010). Moreover, this is specific to the dorsomedial striatum, with no changes in ACh levels in either the dorsolateral striatum or the ventral striatum during reversal learning (Ragozzino *et al.*, 2009).

Despite the importance of striatal ACh in reversal learning, there are currently no studies of striatal ACh function in humans due to a lack of appropriate non-invasive techniques. Proton magnetic resonance spectroscopy (^1H -MRS) is an *in vivo* application of nuclear magnetic resonance spectroscopy that is non-invasive and is implemented on MRI scanners (Puts & Edden, 2012). ^1H -MRS is typically used to measure average metabolite levels in a specific region of interest. However, more recently, ^1H -MRS has been used as a functional measure (functional magnetic resonance spectroscopy, fMRS) to detect event-related fluctuations in brain metabolites during a behavioural task (Apšvalka *et al.*, 2015; Lindner *et al.*, 2017). Although ACh appears on the spectrum at 3.21ppm, several other choline-containing compounds found in the ACh cycle (Figure 1) are also present in the brain at much higher levels, three of which are detectable by MRS; choline (CHO, 3.19ppm), phosphocholine (PC, 3.22ppm) and glycerophosphocholine (GPC, 3.23ppm). Due to its low concentration in the human brain, ACh levels are currently not detectable over the other choline-containing compounds. However, evidence from the animal literature shows that ACh levels are related to CHO levels (Löffelholz, 1998; Klein *et al.*, 2002), making CHO a potential proxy for ACh.

Due to their similarity in chemical structure, CHO, PC and GPC are located close to each other on the MRS spectrum (Figure 2) and are traditionally modelled together as a single peak. This is problematic for the idea of using choline-containing compounds as a measure of ACh (especially when measuring changes in ACh concentration over time), due to the relationship between CHO and PC within the CHO cycle. Indeed, levels of CHO and PC have been shown to be inversely correlated (Miller *et al.*, 1996). Consequently, measuring the three choline-containing compounds as a single summative value will likely mask any functional changes in CHO (e.g. a decrease in CHO would not be observed due to a corresponding increase in PC), making it difficult to capture CHO dynamics that might betray functional ACh concentration fluctuations.

However, we have previously demonstrated that it is possible to separate the CHO peak at 3T whilst modelling PC and GPC together as a single peak (PC+GPC). We further demonstrated task-driven fluctuations of CHO levels that map on to theoretically expected changes in ACh levels, thereby providing support for the potential use of CHO as a proxy measure of ACh dynamics (Lindner *et al.*, 2017).

An important piece of the puzzle then, is the relationship between neuronal activity and CHO levels. Evidence from the animal literature shows a biphasic change in extracellular CHO in response to sustained neuronal activation (Löffelholz, 1998; Klein *et al.*, 2002). There is an initial peak in CHO concentration due to hydrolysis of the released ACh, followed by a decrease below resting levels. This decrease is caused by increased translocation of choline uptake transporters, which mobilise CHO back into the cell and thereby accelerate ACh synthesis. At this point, the rate limiting step of ACh synthesis and release is the availability of free CHO for uptake. Therefore, high temporal resolution in fMRS acquisition would capture both the activity-dependent initial CHO increase (e.g. Linder *et al.* 2017) and later decrease, but lower temporal resolution would only capture the later decrease. Subcortical fMRS acquisitions, such as from the striatum, impose a marked trade-off between temporal resolution and signal-to-noise ratio. As a consequence, here we acquired averaged measurements over short time bins throughout the task, providing a relatively low temporal resolution functional measure, targeting the putative ACh release-dependent decrease in CHO.

Typical studies of human reversal learning use two-choice tasks with multiple reversals. However, during two-choice tasks, once participants have identified the reversal, they straightforwardly switch to the opposite option, guided by their overall model of the task structure. This would provide little information on processes required for learning and expressing a new behaviour. Given evidence that the cholinergic system contributes to new learning after the reversal has been identified, we instead used a multi-alternative probabilistic task, to test the hypothesis that CHO levels in the DS would change during reversal learning. Based on evidence from the animal literature as outlined above, we hypothesised that the sustained involvement of ACh release in reversal learning over multiple task trials would drive a cumulative decrease of free CHO levels.

MATERIALS AND METHODS

Participants

The study was approved by the University of Reading Research Ethics Committee. 45 volunteers (25 female) between the ages of 18.1 and 35.3 (mean = 25.8, SD = 3.8) were recruited by opportunity sampling. All participants were healthy, right handed non-smokers, and gave written informed consent prior to participation.

One participant was excluded due to computer error during data collection. 18 participants were excluded from the analysis reported here as they did not reach the task learning criteria specified below. Of the remaining 26 participants (10 female; mean age = 25.4, SD = 3.2), 15 reached criterion in both the initial learning and reversal rounds (“learners”), and 11 reached criterion in the initial learning round only (“partial learners”).

Behavioural Data

Learning Task

The task was a probabilistic multi-alternative reinforcement learning task, based on the Friedland’s card betting task as adapted for functional magnetic resonance imaging (fMRI) (Schönberg *et al.*, 2007), with an additional reversal component. It was programmed using MATLAB (2014a, The Mathworks, Inc., Natick, MA, United States) and Psychtoolbox (Brainard, 1997).

First, participants were presented with a fixation cross displayed in the centre of the visual display. Participants were then presented with four decks of cards. Each deck contained a mixture of winning and losing cards, corresponding respectively to a gain or loss of 50 points. The probability of getting a winning card differed for each deck (75%, 60%, 40%, and 25%) and the probabilities were randomly assigned across the four decks for each participant. Participants indicated their choice of deck by pressing the corresponding button on a button box. Outcomes were pseudo-randomised so that the assigned probability was true over every 20 times that deck was selected. Additionally, no more than 4 cards of the same result (win/loose) were presented consecutively in the 75% and 25% decks and no more than 3 cards of the same result in the 60% and 40% decks. A cumulative points total was displayed in the bottom right-hand corner throughout the session and in the centre of the visual display at the end of each trial (Figure 3). Participants were instructed that some decks may be better than others, they are free to switch between decks as often as they wish, and they should aim to win as many points as possible.

The learning criterion was set as selection of either of the two highest decks on at least 80% of 20 consecutive trials. As the research question focused on the reversal, we wanted to encourage behavioural and neurochemical stability before the reversal to reduce intra-individual noise. Therefore, a “stability phase” was included at the end of the initial learning phase. The number of trials in this phase was equal to 60% of the number of trials taken to reach criterion. At the end of this phase the deck probabilities were reversed so that the high probability decks became low probability and vice versa. Participants were not informed of the reversal. After reaching the learning criterion again, participants completed a second stability phase, after which the task ended (Figure 4).

Participants were given 100 trials to reach criterion in both the initial learning and reversal phase. If participants did not reach criterion in the initial learning phase, they did not experience the stability phase or the reversal. This was due to the rationale that participants who had not reached criterion during initial learning would not identify a change in contingencies during the reversal and therefore would not behave in the same way as those who had. If participants did not reach criterion in the reversal phase within 100 trials, the task ended and they did not complete the second stability phase.

The presentation timings were jittered. The stimuli were displayed for between 0.8 and 2.8s, with an average display time of 1.7s (standard deviation = 0.6s). Each trial lasted, on average 8.3s (standard deviation = 1.3s). It was essential that the spectral measurements were synchronised with the onset of

the reversal, to ensure there was no overlap of post-criterion stability trials during the measurement acquired over reversal. Therefore, the task was coded so that the onset of the reversal would initiate at the start of a new spectral measurement. If needed, trials were added to the end of the post criterion stability period to ensure the reversal trials occurred during a new spectral measurement (Figure 4).

Control Task

Participants completed a control task identical in structure to that of the learning task, however no feedback was presented. Participants were first presented with a fixation cross in the centre of the visual display. This was followed by the presentation of four decks of cards. Participants indicated their choice of deck using a button box. Participants were then presented with a neutral card, followed by three hashtags in the centre of the screen (i.e. in the location where the overall score would have been presented in the learning task). Participants completed 20 control task trials, corresponding to roughly four minutes. Presentation times were jittered as in the learning task.

Impulsivity

Previous research has shown that trait levels of impulsivity can influence decision making (Bayard *et al.*, 2011). Individuals with higher levels of impulsivity have been shown to demonstrate sub-optimal performance on decision making tasks, displaying a decreased ability to learn reward and punishment associations and implement these to make appropriate decisions. Importantly, individuals with high levels of impulsivity may have difficulty adapting their choice behaviour during reversal learning (Franken *et al.*, 2008; Bayard *et al.*, 2011). Other tasks of cognitive flexibility have also been shown to be influenced by trait impulsivity levels (e.g. Müller *et al.*, 2014). Therefore all participants completed the Barratt Impulsiveness scale (BIS-11; Patton *et al.*, 1995) and their total score was used as a trait measure of impulsivity. This was used in the analysis to verify that any effects of interest were not driven by individual differences in impulsivity and is not reported further.

Performance Variables and Data Analysis

Performance was measured using the number of trials taken to reach criterion in round 1 (initial learning) and in round 2 (which includes the perseveration period and post-reversal learning; Figure 4). ‘Perseveration’ was defined as the number of trials after reversal until the probability of selecting the previously favoured deck reached chance level (0.25), i.e. the number of trials taken to identify the reversal and switch behaviour (this measure was not normally distributed and therefore non-parametric statistical tests were used, e.g. Kendall’s tau-b (t_b) instead of Pearson’s correlation coefficient (r)). ‘Post-reversal learning’ was defined as the number of trials to reach criterion in round 2 minus the number of perseverative trials, i.e. the number of trials to reach the new criterion after the reversal had been detected. In other words, post-reversal leaning was measured by the number of trials the participant took to learn the contingencies once they had realised the deck probabilities had reversed. The number of regressive errors was defined as the number of trials the participant returned to the previously favoured deck after the end of the perseveration period, i.e. after the reversal has been detected.

Magnetic Resonance Spectroscopy

Data Acquisition

Data was collected at the University of Reading on a Siemens Trio 3T MRI scanner using a transmit-receive head coil. A high resolution whole-brain T1 structural image was acquired for voxel placement using a magnetization-prepared, rapid gradient-echo (MPRAGE) sequence parallel to the anterior-posterior commissure line (176 x 1mm slices; TR = 2020ms; TE = 2.9ms; FOV = 250mm).

Voxels were placed in the left dorsal striatum (DS; Figure 5), with anatomy used to guide voxel positioning. The top of the DS was identified by slice-by-slice examination of the structural scan. The slice below the slice where the striatum was no longer visible was selected and the top of the voxel was aligned with this slice. A PRESS sequence was used to acquire data (voxel size = 10mm x 15mm x 15mm; TR = 2000ms; TE = 30ms). 60 spectra were collected and averaged per measurement, creating a two minute ‘spectral bin’. Preliminary tests showed this to be the shortest measurement that could be acquired with adequate signal to noise ratio for quantification. This allowed multiple measurements to be acquired throughout the task, creating a functional measure of metabolite changes. Two initial measurements (average of 60 spectra each) were obtained at rest, followed by a water-suppressed measurement consisting of 15 spectra. Two measurements were acquired during the control task. During the learning task, the number of measurements for each participant varied depending on how long they took to complete the task. Additionally, water-suppressed measurements consisting of 15 spectra were taken before and after both the control and the learning task.

The SIEMENS Auto Align Scout was used to adjust the voxel position based on the actual head position of the participant, correcting for participant motion and minimizing the variability of the voxel position. The scout was used before the structural scan, before the resting acquisition, before the start of the control task, and before the start of the learning task to control for head movement during the scan session.

Structural Segmentation

Structural scans were processed using FSL version 5.0.8 (Smith *et al.*, 2004; Jenkinson *et al.*, 2012). First, the skull was removed using the brain extraction tool (BET; Smith, 2002). Images were segmented into three separate tissue types: grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF) using the FAST tool (Zhang *et al.*, 2001). The coordinates and dimensions of the voxel were then superimposed on these images and the proportion of each of the three tissue types contained within the voxel was calculated.

Quantitation

Data was processed in the time domain using Java-Based Magnetic Resonance User Interface (jMRUI software version 5.0 (<http://www.mrui.uab.es/mrui>); Naressi *et al.*, 2001). Phase correction was performed using the water spectrum acquired closest to the measurement (i.e. the first control spectrum was corrected using the water spectrum acquired before the control task). Each spectrum was then apodized using a Gaussian filter of 3Hz to improve signal quality, reduce noise and reduce effects of signal truncation (Jiru, 2008). The residual water peak was removed using the Hankel-Lanczos Singular Value Decomposition (HLSVD) filter tool.

Metabolite models were generated using the software Versatile Simulation, Pulses and Analysis (VeSPA (<https://scion.duhs.duke.edu/vespa/project>); Soher *et al.*, 2010). 14 typical brain metabolites (Acetate, Aspartate, CHO, Creatine, Gamma-Aminobutyric Acid (GABA), Glucose, Glutamate, Glutamine, Lactate, Myo-inositol, N-acetyl Aspartate (NAA), Phosphocreatine, PC & GPC, Scyllo-inositol, Succinate, Taurine) were simulated at a field strength of 3T using a PRESS pulse sequence (TE1 = 20ms, TE2 = 10ms, main field = 123.25MHz). CHO was modelled separately from PC and GPC (which were modelled as a single peak). Additionally, the sum of the three peaks (total choline, tCHO) was also used in the analyses for comparison with typical practice in choline MRS, and as a control of adequate separation of the CHO in our protocol. As a reminder, we argue that if cholinergic MRS is to be useful as a proxy for ACh function, then it is important to be able to separate CHO from other choline-containing compounds.

Accurate Quantitation of Short Echo time domain signals (AQSES) was applied using the method described in Minati *et al* (2010). To correct for any chemical shift displacement, the spectrum was shifted so that the peak for N-acetyl-aspartate (NAA) was at 2.02ppm. The model was also aligned so that the NAA peak was at 2.02ppm. The frequency range selected for processing was limited to 0-8.6ppm (equal phase for all metabolites, begin time fixed, delta damping (-10 to 25Hz), delta frequency (-5 to 5Hz), no background handling, 0 truncated points, 2048 points in AQSES and normalisation on).

Partial volume correction was performed using the method described by Gasparovic *et al.* (2006) to account for differences in grey matter volume in each voxel. Briefly, the metabolite values were corrected for the concentration of MR visible water based on the proportion of grey matter in the voxel. Additionally, attenuation factors were computed based on T1 and T2 values from grey matter and white matter at 3T for both water and the metabolites. These factors were used to correct the reported values for relaxation effects dependent on the proportion of tissue in the voxel. To control for potential blood-oxygen-level-dependent effects on the signal acquired, and for any differences in signal quality between participants, the corrected values were then referenced against NAA (Zhu & Chen, 2001).

Data Analysis

Only 6% of the data exceeded acceptable Cramer-Rao lower bound (CRLB) thresholds based on criteria described in Posse *et al.*, (2007) (<50% = acceptable reliability). Recent work suggests that data exclusion based on CRLB is not warranted (Kreis, 2016). Nevertheless, data analyses were performed with and without these values. Excluding these values did not alter the results, and therefore results are presented with all values included.

Due to the nature of the task, participants completed different numbers of trials and had variable numbers of spectral bins. Therefore data was aligned at the reversal bin (Rev). The two spectral bins immediately pre-reversal were used for analyses (Rev₋₂, Rev₋₁) as this was the maximum number of bins preceding the reversal that contained data for all participants. Additionally, the two spectral bins immediately post-reversal were chosen for consistency (Rev₊₁, Rev₊₂) (all participants except one had data for both post-reversal bins). To control for individual differences in average metabolite levels, the metabolite levels for each bin are reported as a change from the participant's average metabolite level over each task.

Statistical analysis was performed using SPSS (IBM Corp. Released 2013. IBM SPSS Statistics for Windows, Version 22.0. Armonk, NY: IBM Corp).

Confounding Variables

To check that variations in metabolite values were not due to varying tissue composition in the voxels of different participants, we tested the correlation between cholinergic metabolite levels (average metabolite levels at rest, during the control period, and during the task period) and proportion of grey matter, white matter and the water signal. No significant correlations were found and therefore these are not reported further.

There is evidence that metabolite levels in the brain can vary based on time of day (Soreni *et al.*, 2006) and age (Pfefferbaum *et al.*, 1999; Reyngoudt *et al.*, 2012). To check for any such effects, we tested the correlation between these factors and the average metabolite levels during the control task and during the learning task. No significant correlations were found and therefore these are not reported further.

RESULTS

Task Performance

Twenty six (26) participants reached our pre-defined learning criterion and progressed to the reversal learning phase. Fifteen (15) of these participants went on to achieve the same accuracy criterion again after reversal (we refer to these participants as learners; the remaining 11 participants never reached criterion after reversal, and we refer to them as partial learners). Participant performance is summarised in Table 1.

Resting Average CHO Levels

Learners (i.e. participants who went on to reach the learning criterion in both rounds; N=15) had significantly lower average levels of DS CHO (but not PC+GPC) at rest, compared to partial learners (i.e. participants who only reached criterion during the initial learning phase, but not after reversal learning; N=11). This effect was independent of the participants' initial learning efficiency as indexed by numbers of trials-to-criterion (between group ANCOVA, controlling for initial learning (R1) trials-to-criterion: $F(1)=9.043, p=0.006$, partial $\eta^2=0.282$). Individual rest levels of CHO did not alter the task-related changes reported in the following section.

Task-related Functional Changes in CHO Levels

During the task, spectral acquisition bins were aligned to the start of testing (initial learning) and the onset of the first reversal bin. Paired-samples t-tests were conducted over consecutive bins during the initial learning and reversal periods for all participants who went through reversal (N=26) to assess changes in CHO levels.

In line with our prediction, during reversal learning there was a specific significant drop in CHO levels during the first reversal bin compared to the bin taken directly preceding the reversal (which corresponds to the imposed stability period of the task; see Methods for more details) (Table 2). This drop in CHO levels mirrored the large drop in accuracy levels during the first reversal bin, reflecting the unexpected change in contingencies (Figure 6). This drop was specific to reversal; no CHO changes were observed at any other point during either the learning task or the control task (Figure 6, Table 2). Individual differences in CHO concentration changes during the first reversal bin were not associated with the efficiency of the reversal (as indexed by the number of perseverative responses ($t_b(23)=0.021, p=0.894$), or with the number of regressive errors ($r(23) = -0.304, p = 0.159$), and there

was no difference between learners and partial learners in the magnitude of CHO change during the first reversal bin ($t(24)=-0.684$, $p=0.500$).

To test the specificity of the CHO effect reported above, we also checked for changes in the PC+PGC and tCHO estimates, as well as control metabolites total creatine (tCre; sum of creatine and phosphocreatine) and glutamate (Glu). tCre is thought to be stable in brain tissue (Rae, 2014), and was not expected to fluctuate with task epoch. Glu provides an extra functional control given that any Glu changes (potentially associated with neurotransmitter recruitment during the task) would not show the specificity or direction of the predicted CHO changes (which were shown to be concentration reductions specifically during the reversal period). None of these tests were significant (summarised in Figure 6 and Table 2), and therefore support the specificity of the reported effect of interest.

DISCUSSION

We used fMRS during a probabilistic reversal learning task to provide direct evidence of the involvement of the human striatal cholinergic system in reversal learning. We found a task-related decrease in CHO levels specific to reversal, but not initial, learning. Based on the dynamics of the ACh cycle, **we propose that this may reflect a sustained increase in ACh release**, analogous to that found in animal studies (Ragozzino *et al.*, 2009; Brown *et al.*, 2010).

During tasks that impose a change in action-outcome contingencies, evidence from the animal literature shows a sustained increase in ACh levels throughout the reversal learning period. Here we observed a drop in CHO levels during reversal, in line with a hypothesis derived from animal experiments showing a drop on CHO concentration following sustained activity-dependent ACh release (Klein *et al.*, 2002).

The primary target of ACh in the striatum is medium spiny neurons (MSNs), which make up the majority of the striatum. MSNs integrate multiple inputs from the cortex to assist in action selection and goal directed behaviour. By interacting with MSNs, ACh is thought to provide context-sensitive learning, for example, when an action may lead to a reward in one context, but not in another (Ashby & Crossley, 2011). ACh has an inhibitory effect on MSN activity, and it has been hypothesised that changes in ACh levels may be used to select the appropriate set of MSNs for association-outcome encoding (Stalnaker *et al.*, 2016). Indeed, a recent study by Stalnaker *et al.* (2016) has demonstrated the crucial role of dynamic CIN activity (the main source of ACh in the striatum) in selecting context-appropriate actions. Rats were trained on a 3-choice behavioural task which required alternate responses depending on the trial context, indicated by the presence of an odour. Depending on the odour presented, the rat would receive a reward in one of three fluid wells. If the rat went to the wrong fluid well, they would not receive a reward. Not only did CINs signal the corresponding state of the trial (i.e. which fluid well would be rewarded), they also showed that, on trials where the rats did not make the correct or optimal choice, the block decoding by the CINs was not accurate i.e. the CINs were not accurately identifying the correct trial state. Moreover, this effect was specific to CIN activity, and was not seen in MSNs. This demonstrates the importance of CIN function in specifying the action relevant context, without which inappropriate/suboptimal behaviour occurs (Stalnaker *et al.*, 2016).

Context signalling is an important aspect for reversal learning and cognitive flexibility more generally. After a change in context or contingencies, the previous action-outcome association is no longer reinforced. Therefore, without context signalling, this association would decay and would need to be re-learnt when the original context was experienced again. However, the inhibitory effect of

ACh prevents the decay of the original action-outcome association. In this way, the existing action-outcome association is protected and a new action-outcome association can be learnt, without removal of the existing knowledge. This is supported by evidence from Bradfield *et al.* (2013), who showed that disruption of CIN signalling in rats through disruption of thalamic inputs resulted not only in an inability to encode a change in contingencies, but also an inability to express the existing learning.

In our study, following the initial decrease, CHO levels progressively increased during post-reversal learning. Evidence in rodents shows that CHO levels eventually return to baseline during sustained CIN firing, even though ACh levels may remain elevated (Klein *et al.*, 2002). Therefore, based on our data alone we cannot infer whether CHO levels returned to baseline due to an activity-linked reduction in ACh release (i.e. reduction in CIN firing), or if they simply returned to equilibrium **as suggested by the rodent evidence**. Consequently we cannot deduce whether ACh levels remain elevated throughout post-reversal learning, or indeed if further, non-linear changes in firing rate occur. Additionally, evidence regarding the profile of ACh release in rodents during reversal is conflicting. An initial study by Ragazzino & Choi (2004) showed that ACh levels returned to baseline during post-reversal learning, while a later study by Ragazzino *et al* (2009) showed that the ACh increase was sustained throughout post-reversal learning. Further work is needed to determine the dynamics of ACh and CHO release and recovery post-reversal.

There was no association between CHO changes and performance during reversal. There are a number of reasons for this. Firstly, in order to reduce functional heterogeneity, our sample reached a stringent performance criterion which they maintained for a protracted stability period before the reversal was implemented. Measureable individual differences in performance are consequently reduced.

Secondly, the CHO measure has both a low spatial and temporal resolution. Animal studies have shown that the role of the cholinergic system during reversal is specific to the dorso-medial striatum (roughly analogous to the human associative caudate), with no effects seen in the dorso-lateral striatum (roughly analogous to the human posterior putamen). Though there have been no studies investigating spatial specificity of the cholinergic system in humans, Bernácer *et al.* (2007) demonstrated that the density of CINs is roughly 1.8 times higher in the caudate than the putamen, suggesting that there will be a similar spatial specificity of activity. The size of the voxel needed in this study forces a currently inevitable compromise of spatial specificity. Additionally, due to the low temporal resolution and the fact that participants learnt at different rates, the spectral bin acquired at reversal will have contained trials from different stages of learning. Therefore, any changes in CHO levels related to individual differences in learning will be masked in this measurement.

Finally, the prominent effect following disruption of cholinergic signalling in animals is an increase in the number of regressive errors (returning to the previously relevant behaviour after a change in behaviour has occurred) as opposed to perseverative errors (continued selection of the previously rewarded stimulus) (Bradfield *et al.*, 2013). However, this aspect of animal behaviour cannot be directly compared with the behaviour of humans in multi-alternative choice tasks. Human participants will base their decisions in part on a higher-level representation of the task (potentially a verbalisable rule, e.g. “the left-most deck is now the most rewarded”) which is thought to be maintained in frontal areas of the cortex (Armbruster *et al.*, 2012). Therefore the number of regressive errors will likely be modulated by task representation, with participants with the most stable task representation making less regressive errors. Recent evidence from the animal literature shows that, although rats with disrupted striatal cholinergic signalling were initially impaired during reversal, they were able to learn and express the reversed contingencies after additional training (Bradfield & Balleine, 2017). It is

possible that this occurs on a faster timescale with human participants. It should be noted that the cortico-basal ganglia-thalamo-cortical (CBGTC) system has been shown to be modulated by the maintenance of task rules, with those with stronger representation of the task structure showing higher activation in the caudate and thalamus (the main input to the CINs) during a behaviour switch (Ueltzhöffer *et al.*, 2015). These findings suggest that cortical task structure representation likely “bootstraps” the involvement of DS activity in behavioural flexibility (Paul & Ashby, 2013), possibly via thalamostriatal input to the CINs. This highlights the need for further study of this system, to bridge the gap between our understanding of CIN/cholinergic dynamics in animals and humans, and to integrate the role of the striatal cholinergic system in the better understood dopaminergic modulation of CBGTC function.

To test the specificity of the change in CHO levels during reversal, we compared it to changes in CHO levels during a simple button-pressing control task, identical in structure to the learning task but with no learning component. There were no significant changes in CHO levels during the control task, demonstrating that the decrease is task related and not due to a non-specific fluctuation in CHO levels. Additionally, there were no significant changes in CHO levels between the other bins obtained during the task, demonstrating that this decrease is related to the reversal only and not the task in general. Further, the observed changes were specific to CHO, and were not seen in PC+GPC, or when all three choline-containing compounds were modelled together (tCHO), emphasising the need to model CHO separately when investigating CHO levels as a proxy for ACh levels, as demonstrated in Lindner *et al.* (2017). Additionally, no significant changes were seen in the control metabolites tCR and Glu, suggesting that the result is task specific and not related to the measurement procedure itself, or general fluctuation in metabolite levels.

The application of MRS to the study of the cholinergic system is undoubtedly in its infancy. Nevertheless, the results of this study, in conjunction with our previous work demonstrating task-dependent fluctuations in cortical CHO during visuospatial attention shifts (Lindner *et al.*, 2017), are promising in this regard. Notably, both studies tackle the more ambitious element of the cholinergic MRS agenda, namely the effort to indirectly capture, through CHO tracking, the dynamics of acetylcholine function.

These promising findings notwithstanding, it is important to emphasise that both our studies demonstrate function-relevant changes in CHO levels, but provide no evidence regarding the relationship of these changes to ACh release. Numerous questions and technological challenges remain before cholinergic MRS can be considered a mainstream methodology.

Although beyond the scope of this paper, a key challenge for this work is the lack of data on the relative concentration of choline-containing metabolites in different parts of the brain.

It has been shown that GPC, PC and CHO are present in approximate ratios of 1:0.3:0.2 (Ala-Korpela *et al.*, 1996). However, these ratios are based on measurements acquired in the hippocampus, which receives its cholinergic supply from the basal forebrain. In the case of the striatum, the local cholinergic system is unique, not only in containing by far the highest concentration of ACh in the brain (Hoover *et al.*, 1978), but importantly in that the majority of striatal ACh is provided via an interneuronal system (Aosaki *et al.*, 2010). In addition, these CINs are tonically active (Goldberg & Reynolds, 2011), resulting in an ambient level of ACh in the striatum. Therefore, levels of CHO are likely to be much higher in the striatum compared to other areas of the brain. However, to the authors' knowledge, there has been no investigation into the relative levels of GPC, PC and CHO in the striatum.

In animal studies, it is possible to manipulate and measure the action of ACh directly, and consequently there is little interest in the measurement of the magnitude and dynamics of activity-related CHO concentration. Direct measurement of CHO in animal models (especially regional CHO concentration changes and diffusion characteristics in response to neural activation, and their relationship to activity-dependent ACh release) would significantly increase our confidence in the human measurements, and could potentially allow us to build a precise mathematical model of the functional relationship between ACh and CHO (as well as other CHO-containing metabolites).

Further, we note in particular the difficulties posed by the very large voxel size needed at 3T to acquire an acceptable signal-to-noise ratio, which is problematic for anatomically specific questions. We also highlight the related issue of poor temporal resolution, advances in which would provide significant impetus for further development of the method.

In summary, using fMRS, and for the first time in humans, we have shown a task-related decrease in CHO levels in the human DS during reversal learning. We interpret this to reflect a sustained post-reversal increase in ACh levels, which is in line with findings from the animal literature. **Further refinement of this method is needed, in particular with reference to time-resolved subcortical MRS measurements at 3T, and evidence from invasive animal models regarding the activity-dependent CHO-ACh relationship.** Nevertheless, this novel finding provides further support for the idea of developing event-related ^1H -MRS for the *in vivo* study of the human cholinergic system.

ACKNOWLEDGMENTS

This study was supported by a Human Frontier Science Program grant (RGP0048/2012) and an Engineering and Physical Sciences Research Council doctoral training grant (EP/L505043/1). The funders had no involvement in study design, the collection, analysis, and interpretation of data; in writing the report, or in the decision to submit the paper for publication. The authors would like to thank Dr Jayne Morriss, Nadia ElHorch and Raul Ungureanu for assistance with data collection.

COMPETING INTERESTS

The authors declare no conflicts of interest.

AUTHOR CONTRIBUTIONS

Conceptualization: AC

Data curation: TB AC

Formal analysis: TB ML PGM AC

Funding acquisition: AC

Investigation: TB

Methodology: TB ML PGM AC

Project administration: AC

Resources: PGM AC

Software: TB ML PGM

Supervision: AC

Validation: TB ML PGM AC

Visualization: TB AC

Writing – original draft: TB

Writing – review & editing: TB ML PGM AC

DATA ACCESSIBILITY STATEMENT

Anonymised data are freely available on the Open Science Framework: Bell, T.K., Lindner, M., Mullins, P.G., & Christakou, A. (2017) Functional neurochemical imaging of the human striatal cholinergic system during reversal learning [WWW Document]. URL <https://osf.io/hd5rj>.

ABBREVIATIONS

accurate quantitation of short echo time domain signals (AQSES)

acetyl coenzyme A (acetyl-CoA)

acetylcholine (ACh)

acetylcholinesterase (AChE)

alkaline phosphatase (AP)

choline (CHO)

choline acetyltransferase (ChAT)

choline kinase (CK)

cholinergic interneuron (CIN)

cortico-basal ganglia-thalamo-cortical (CBGTC)

Cramer-Rao lower bound (CRLB)

creatine (CRE)

cytidine diphosphocholine (CDP-CHO)
dorsal striatum (DS)
field of view (FOV)
functional-¹H-MRS (fMRS)
gamma-aminobutyric acid (GABA)
glutamate (Glu)
glycerophosphocholine (GPC)
GPC hydrolase (GPC-H)
Kendall's tau-b (t_b)
lyso-phosphatidylcholine (Lyso-PtdCHO)
lysophospholipase (LPL)
magnetization-prepared, rapid gradient-echo (MPRAGE)
Montreal Neurological Institute (MNI)
N-acetyl aspartate (NAA)
PC cytidyltransferase (PC-CT)
Pearson's correlation coefficient (r)
phosphatidylcholine (PtdCHO)
phosphocholine (PC)
phospholipase (PL)
phosphocreatine (PCRE)
point-resolved spectroscopy (PRESS)
proton magnetic resonance spectroscopy (¹H-MRS)
standard deviation (SD)
time-to-echo (TE)
time-to-repetition (TR)
total choline (tCHO)
total creatine (tCre)
versatile simulation, pulses and analysis (VeSPA)

REFERENCES

Ala-Korpela, M., Posio, P., Mattila, S., Korhonen, A., & Williams, S.R. (1996) Absolute quantification of phospholipid metabolites in brain-tissue extracts by ^1H NMR spectroscopy. *J. Magn. Reson. B*, **113**, 184–189.

Aosaki, T., Miura, M., Suzuki, T., Nishimura, K., & Masuda, M. (2010) Acetylcholine-dopamine balance hypothesis in the striatum: An update. *Geriatr. Gerontol. Int.*, **10**, S148–S157.

Apšvalka, D., Gadie, A., Clemence, M., & Mullins, P.G. (2015) Event-related dynamics of glutamate and BOLD effects measured using functional magnetic resonance spectroscopy (fMRS) at 3 T in a repetition suppression paradigm. *Neuroimage*, **118**, 292–300.

Armbruster, D.J.N., Ueltzhöffer, K., Basten, U., & Fiebach, C.J. (2012) Prefrontal cortical mechanisms underlying individual differences in cognitive flexibility and stability. *J. Cogn. Neurosci.*, **24**, 2385–2399.

Ashby, F.G. & Crossley, M.J. (2011) A computational model of how cholinergic interneurons protect striatal-dependent learning. *J. Cogn. Neurosci.*, **23**, 1549–1566.

Bayard, S., Raffard, S., & Gely-Nargeot, M.-C. (2011) Do facets of self-reported impulsivity predict decision-making under ambiguity and risk? Evidence from a community sample. *Psychiatry Res.*, **190**, 322–326.

Bernácer, J., Prensa, L., & Giménez-Amaya, J.M. (2007) Cholinergic interneurons are differentially distributed in the human striatum. *PLoS One*, **2**, e1174.

Boulanger, Y., Labelle, M., & Khiat, A. (2000) Role of phospholipase A2 on the variations of the choline signal intensity observed by ^1H magnetic resonance spectroscopy in. *Brain Res. Rev.*, **33**, 380–389.

Bradfield, L.A. & Balleine, B.W. (2017) Thalamic control of dorsomedial striatum regulates internal state to guide goal-directed action selection. *J. Neurosci.*, **37**, 3860–16.

Bradfield, L.A., Bertran-Gonzalez, J., Chieng, B., & Balleine, B.W. (2013) The thalamostriatal pathway and cholinergic control of goal-directed action: interlacing new with existing learning in the striatum. *Neuron*, **79**, 153–166.

Brainard, D.H. (1997) The Psychophysics Toolbox. *Spat. Vis.*, **10**, 433–436.

Brown, H.D., Baker, P.M., & Ragazzino, M.E. (2010) The parafascicular thalamic nucleus concomitantly influences behavioral flexibility and dorsomedial striatal acetylcholine output in rats. *J. Neurosci.*, **30**, 14390–14398.

Franken, I.H.A., van Strien, J.W., Nijs, I., & Muris, P. (2008) Impulsivity is associated with behavioral decision-making deficits. *Psychiatry Res.*, **158**, 155–163.

Gasparovic, C., Song, T., Devier, D., Bockholt, H.J., Caprihan, A., Mullins, P.G., Posse, S., Jung, R.E., & Morrison, L.A. (2006) Use of tissue water as a concentration reference for proton spectroscopic imaging. *Magn. Reson. Med.*, **55**, 1219–1226.

Goldberg, J. a & Reynolds, J.N.J. (2011) Spontaneous firing and evoked pauses in the tonically active cholinergic interneurons of the striatum. *Neuroscience*, **198**, 27–43.

Hoover, D.B., Muth, E. a, & Jacobowitz, D.M. (1978) A mapping of the distribution of acetylcholine, choline acetyltransferase and acetylcholinesterase in discrete areas of rat brain. *Brain Res.*, **153**, 295–306.

Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., & Smith, S.M. (2012) FSL. *Neuroimage*, 782–790.

Jiru, F. (2008) Introduction to post-processing techniques. *Eur. J. Radiol.*, **67**, 202–217.

Klein, J., Weichel, O., Ruhr, J., Dvorak, C., & Loffelholz, K. (2002) A homeostatic mechanism counteracting K+-evoked choline release in adult brain. *J. Neurochem.*, **80**, 843–849.

Kreis, R. (2016) The trouble with quality filtering based on relative Cramér-Rao lower bounds. *Magn. Reson. Med.*, **75**, 15–18.

Lindner, M., Bell, T., Iqbal, S., Mullins, P.G., & Christakou, A. (2017) In vivo functional neurochemistry of human cortical cholinergic function during visuospatial attention. *PLoS One*, **12**, e0171338.

Lockman, P.R. & Allen, D.D. (2002) The transport of choline. *Drug Dev. Ind. Pharm.*, **28**, 749–771.

Löffelholz, K. (1998) Brain choline has a typical precursor profile. *J. Physiol.*, **92**, 235–239.

McCool, M.F., Patel, S., Talati, R., & Ragozzino, M.E. (2008) Differential involvement of M1-type and M4-type muscarinic cholinergic receptors in the dorsomedial striatum in task switching. *Neurobiol. Learn. Mem.*, **89**, 114–124.

Miller, B.L., Changl, L., Booth, R., Ernst, T., D, M., Ph, D., Nikas, D., McBride, D., & Jenden, D.J. (1996) In vivo 1H-MRS Choline: Correlation with in vitro chemistry/histology. *Life Sci.*, **58**, 1929–1935.

Minati, L., Aquino, D., Bruzzone, M., & Erbetta, A. (2010) Quantitation of normal metabolite concentrations in six brain regions by in-vivo 1 H-MR spectroscopy. *J. Med. Phys.*, **35**, 154.

Müller, V.I., Langner, R., Cieslik, E.C., Rottschy, C., & Eickhoff, S.B. (2015) Interindividual differences in cognitive flexibility: influence of gray matter volume, functional connectivity and trait impulsivity. *Brain Struct. Funct.*, **220**, 2401–2414.

Naressi, A., Couturier, C., Devos, J.M., Janssen, M., Mangeat, C., de Beer, R., & Graveron-Demilly, D. (2001) Java-based graphical user interface for the MRUI quantitation package. *Magn. Reson. Mater. Physics, Biol. Med.*, **12**, 141–152.

Patton, J.H., Stanford, M.S., & Barratt, E.S. (1995) Factor structure of the barratt impulsiveness scale. *J. Clin. Psychol.*, **51**, 768–774.

Paul, E.J. & Ashby, F.G. (2013) A neurocomputational theory of how explicit learning bootstraps early procedural learning. *Front. Comput. Neurosci.*, **7**, 177.

Pfefferbaum, A., Adalsteinsson, E., Spielman, D., Sullivan, E. V., & Lim, K.O. (1999) In vivo

spectroscopic quantification of the N-acetyl moiety, creatine, and choline from large volumes of brain gray and white matter: Effects of normal aging. *Magn. Reson. Med.*, **41**, 276–284.

Posse, S., Otazo, R., Caprihan, A., Bustillo, J., Chen, H., Henry, P.-G., Marjanska, M., Gasparovic, C., Zuo, C., Magnotta, V., Mueller, B., Mullins, P., Renshaw, P., Ugurbil, K., Lim, K.O., & Alger, J.R. (2007) Proton echo-planar spectroscopic imaging of J-coupled resonances in human brain at 3 and 4 Tesla. *Magn. Reson. Med.*, **58**, 236–244.

Puts, N.A.J. & Edden, R.A.E. (2012) In vivo magnetic resonance spectroscopy of GABA: A methodological review. *Prog. Nucl. Magn. Reson. Spectrosc.*, **60**, 29–41.

Rae, C.D. (2014) A guide to the metabolic pathways and function of metabolites observed in human brain ^1H magnetic resonance spectra. *Neurochem. Res.*, **39**, 1–36.

Ragozzino, M.E. & Choi, D. (2004) Dynamic changes in acetylcholine output in the medial striatum during place reversal learning. *Learn. Mem.*, **11**, 70–77.

Ragozzino, M.E., Jih, J., & Tzavos, A. (2002) Involvement of the dorsomedial striatum in behavioral flexibility: role of muscarinic cholinergic receptors. *Brain Res.*, **953**, 205–214.

Ragozzino, M.E., Mohler, E.G., Prior, M., Palencia, C.A., & Rozman, S. (2009) Acetylcholine activity in selective striatal regions supports behavioral flexibility. *Neurobiol. Learn. Mem.*, **91**, 13–22.

Reyngoudt, H., Claeys, T., Vlerick, L., Verleden, S., Acou, M., Deblaere, K., De Deene, Y., Audenaert, K., Goethals, I., & Achten, E. (2012) Age-related differences in metabolites in the posterior cingulate cortex and hippocampus of normal ageing brain: a ^1H -MRS study. *Eur. J. Radiol.*, **81**, e223-31.

Schönberg, T., Daw, N.D., Joel, D., & O'Doherty, J.P. (2007) Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *J. Neurosci.*, **27**, 12860–71286.

Smith, S.M. (2002) Fast robust automated brain extraction. *Hum. Brain Mapp.*, **17**, 143–155.

Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobniak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., & Matthews, P.M. (2004) Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, **23**, S208–S219.

Soher, B.J., Semanchuk, D., Todd, D., Steinberg, J., & Young, K. (2011) Vespa: Integrated applications for RF pulse design, spectral simulation and MRS data analysis. *Proc. 19th Annu. Meet. ISMRM*, **19**, 1410.

Soreni, N., Noseworthy, M.D., Cormier, T., Oakden, W.K., Bells, S., & Schachar, R. (2006) Intraindividual variability of striatal (1)H-MRS brain metabolite measurements at 3 T. *Magn. Reson. Imaging*, **24**, 187–194.

Stalnaker, T.A., Berg, B., Aujla, N., & Schoenbaum, G. (2016) Cholinergic interneurons use orbitofrontal input to track beliefs about current state. *J. Neurosci.*, **36**, 6242–6257.

Tzavos, A., Jih, J., & Ragozzino, M.E. (2004) Differential effects of M1 muscarinic receptor blockade and nicotinic receptor blockade in the dorsomedial striatum on response reversal learning. *Behav. Brain Res.*, **154**, 245–253.

Ueltzhöffer, K., Armbruster-Genç, D.J.N., & Fiebach, C.J. (2015) Stochastic dynamics underlying cognitive stability and flexibility. *PLoS Comput. Biol.*, **11**, e1004331.

Zhang, Y., Brady, M., & Smith, S.M. (2001) Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging*, **20**, 45–57.

Zhu, X.-H. & Chen, W. (2001) Observed BOLD effects on cerebral metabolite resonances in human visual cortex during visual stimulation: A functional ¹H MRS study at 4 T. *Magn. Reson. Med.*, **46**, 841–847.

TABLES

Table 1. Number of trials per task phase

	N	Average Number of Trials	SD
Initial Learning (to criterion)	26	49	24
First Stability Phase	26	32	14
Reversal Learning			
Perseveration Period	23	14	9
Number of Regressive Errors	23	11	8
Post Reversal Learning (to criterion)	15	39	20
Second Stability Phase	15	34	11
Total	15	149	46

Twenty six (26) participants reached criterion during the initial learning phase of the task and proceeded to the reversal learning. Fifteen (15) of these participants reached the learning criterion following the reversal (the task was terminated after 100 post-reversal trials if criterion was not reached; N=11).

Table 2. Paired samples *t*-tests comparing accuracy and metabolite concentration changes across spectrum acquisition bins

	Accuracy				CHO				PC+GPC				tCHO				tCre				Glu			
	t	df	p	p _B	t	df	p	p _B	t	df	p	p _B	t	df	p	p _B	t	df	p	p _B	t	df	p	p _B
Control Task																								
CT1-CT2	N/A	N/A	N/A	N/A	-1.812	25	0.082	0.085	-0.020	25	0.985	0.983	-1.862	25	0.074	0.074	0.348	25	0.731	0.732	-0.665	25	0.512	0.513
Initial Learning																								
IL1-IL2	-2.397	25	0.024	0.025	-0.836	24	0.412	0.416	1.030	24	0.313	0.317	0.377	24	0.709	0.716	1.848	24	0.770	0.120	0.533	24	0.585	0.594
IL2-IL3	-6.229	21	0.000	0.000	0.440	20	0.664	0.674	-1.244	20	0.228	0.231	-1.377	20	0.184	0.183	-0.197	20	0.769	0.774	-1.125	20	0.274	0.269
Reversal Learning																								
Rev ₋ -Rev ₁	-2.035	25	0.053	0.056	-0.027	25	0.978	0.980	1.490	25	0.149	0.157	0.726	25	0.474	0.476	1.519	24	0.142	0.147	0.128	25	0.899	0.899
Rev ₁ -Rev	11.049	25	0.000	0.000	2.412	25	0.024	0.022	-0.952	25	0.350	0.347	1.183	25	0.248	0.250	-0.496	24	0.624	0.621	0.548	25	0.589	0.599
Rev-Rev ₊₁	-5.082	25	0.000	0.000	-0.728	25	0.473	0.466	-1.049	25	0.304	0.307	-1.299	25	0.206	0.206	0.882	25	0.386	0.380	-0.752	25	0.459	0.456
Rev ₊₁ -Rev ₊₂	-2.405	23	0.025	0.031	-1.330	24	0.196	0.193	1.385	24	0.179	0.180	-0.020	24	0.985	0.984	0.284	24	0.779	0.775	-0.616	24	0.543	0.540

Note: C_n denotes the control task bins; IL_n denotes the initial learning bins. Rev_n denotes the position of the spectral bin in relation to the reversal bin (Rev); CHO = choline; PC+GPC = phosphocholine and glycerophosphocholine (modelled as one peak); tCHO = total choline; tCre = total creatine; Glu = glutamate. Numbers in bold denote statistically significant comparisons. “p_B” refers to the significance test of a 10,000 sample bootstrap analysis at 95% confidence interval.

FIGURE CAPTIONS

Figure 1: The CHO cycle

Phosphatidylcholine (PtdCHO) is stored in the cell membrane. When needed, it is retrieved and converted to lyso-phosphatidylcholine (Lyso-PtdCHO) by phospholipase (PL). Lyso-PtdCHO is then transformed to glycerophosphocholine (GPC) by lysophospholipase (LPL). The enzyme GPC hydrolase (GPC-H) hydrolyses GPC to phosphocholine (PC). PC is then converted to choline (CHO) by alkaline phosphatase (AP). Acetylcholine (ACh) is synthesised from CHO and acetyl coenzyme A (acetyl-CoA) by choline acetyltransferase (ChAT). ACh is then released into the synapse and broken down into CHO and acetate by acetylcholinesterase (AChE). CHO is taken back up into the synapse by choline uptake transporters, where it is converted back into PC by choline kinase (CK). PC is then converted to cytidine diphosphocholine (CDP-CHO) by the enzyme PC cytidyltransferase (PC-CT), which is then converted to PtdCHO by diacylglycerolholine phosphotransferase for storage back in the membrane (Boulanger, Labelle, & Khiat, 2000; Lockman & Allen, 2002) Figure adapted from Lindner et al., 2017.

Figure 2: Spectrum of the choline-containing compound peaks

Locations of the spectral metabolite peaks for choline (CHO), phosphocholine (PC) and glycerophosphocholine (GPC), taken from Lindner et al., 2017. The three peaks are located close together on the spectrum, and are often modelled together as a single peak. The choline-containing compounds were simulated using VeSPA (<https://scion.duhs.duke.edu/vespa/project>) at a field strength of 3T (main field 123.25MHz) with a PRESS pulse sequence (TE1 = 20ms, TE2 = 10ms).

Figure 3: Task trial schematic

Participants were instructed to choose between four decks of cards. Each deck had a different probability of generating winning cards (75%, 60%, 40% and 25%). Once the learning criterion had been reached, the deck probabilities were flipped so that high probability decks became low probability decks and vice versa. Participants were not informed of this in advance and were simply instructed to gain as many points as possible. RT = reaction time.

Figure 4: General overview of learning task structure

Upon reaching criterion in the initial learning phase (round 1 or R1), participants then completed a post criterion stability phase (criterion 1 or C1), equal in number to 60% of the number of trials taken to reach criterion. After this phase, the deck probabilities were reversed. ^1H -MRS onset was time locked to the onset of reversal. If needed, trials were added to the end of the post criterion stability period to ensure the reversal trials occurred during a new spectral measurement. Participants then completed a reversal learning phase (round 2 or R2) and upon reaching criterion again, they completed a post criterion stability phase (criterion 2 or C2), again equal in number to 60% of the number of trials taken to reach criterion. The number of trials completed in each stage varied depending on each participant's performance on the task.

Figure 5: Sum of voxel placement over all subjects and representative single spectrum

Left panel: The ^1H -MRS voxel masks were transformed into MNI space for better visualisation. All voxels encompassed the left dorsal striatum (centre MNI coordinates: x = -22.48, y = -5.61, z =

11.94). Right panel: Example of spectral data after pre-processing from a single spectral bin acquired during the learning task from one participant. (“data”). Included are the model peaks for choline (CHO), creatine (CRE), phosphocholine and glycerophosphocholine (PC&GPC), n-acetyl aspartate (NAA), phosphocreatine (PCRE) and the residual. –

Figure 6: Changes in metabolite levels during the initial learning and reversal, superimposed on accuracy levels

Top panel: CHO concentration change during initial learning bins (ILn) and reversal (Rev-n: pre-reversal bins, Rev: reversal bin (indicated by an arrow), Rev+n: post-reversal bins), superimposed over group mean accuracy levels. A significant decrease in CHO levels was observed specifically during the reversal bin (Rev) (the asterisk denotes a significant drop in both the CHO and accuracy levels; $p<0.05$). No other changes in CHO levels during either the control task or the learning task were statistically significant. Bottom panel: Choline-containing compounds fluctuations during initial learning and reversal, superimposed over group mean accuracy levels. No significant changes were observed in combined levels of PC and GPC (PC+GPC; modelled together as a single peak), or when modelling all three choline containing compounds as a single peak (tCHO). For further details, and details of comparisons not shown here, please see Table 2. Each bin is defined by a single spectral acquisition. Error bars denote the standard error.







