



Local Fields in Human Subthalamic Nucleus Track the Lead-up to Impulsive Choices

John M. Pearson^{1*}, Patrick T. Hickey², Shivanand P. Lad³, Michael L. Platt^{1,4} and Dennis A. Turner^{3,4}

¹ Center for Cognitive Neuroscience and Duke Institute for Brain Sciences, Duke University, Durham, NC, United States,

² Department of Neurology, School of Medicine, Duke University, Durham, NC, United States, ³ Department of Neurosurgery, School of Medicine, Duke University, Durham, NC, United States, ⁴ Department of Neurobiology, School of Medicine, Duke University, Durham, NC, United States

OPEN ACCESS

Edited by:

O'Dhanial A. Mullette-Gillman,
National University of Singapore,
Singapore

Reviewed by:

Claudio Lucchiari,
Università degli Studi di Milano, Italy
Wael Asaad,
Brown University, United States

*Correspondence:

John M. Pearson
john.pearson@duke.edu

Specialty section:

This article was submitted to
Decision Neuroscience,
a section of the journal
Frontiers in Neuroscience

Received: 09 August 2017

Accepted: 07 November 2017

Published: 23 November 2017

Citation:

Pearson JM, Hickey PT, Lad SP,
Platt ML and Turner DA (2017) Local
Fields in Human Subthalamic Nucleus
Track the Lead-up to Impulsive
Choices. *Front. Neurosci.* 11:646.
doi: 10.3389/fnins.2017.00646

The ability to adaptively minimize not only motor but cognitive symptoms of neurological diseases, such as Parkinson's Disease (PD) and obsessive-compulsive disorder (OCD), is a primary goal of next-generation deep brain stimulation (DBS) devices. On the basis of studies demonstrating a link between beta-band synchronization and severity of motor symptoms in PD, the minimization of beta band activity has been proposed as a potential training target for closed-loop DBS. At present, no comparable signal is known for the impulsive side effects of PD, though multiple studies have implicated theta band activity within the subthalamic nucleus (STN), the site of DBS treatment, in processes of conflict monitoring and countermanding. Here, we address this challenge by recording from multiple independent channels within the STN in a self-paced decision task to test whether these signals carry information sufficient to predict stopping behavior on a trial-by-trial basis. As in previous studies, we found that local field potentials (LFPs) exhibited modulations preceding self-initiated movements, with power ramping across multiple frequencies during the deliberation period. In addition, signals showed phasic changes in power around the time of decision. However, a prospective model that attempted to use these signals to predict decision times showed effects of risk level did not improve with the addition of LFPs as regressors. These findings suggest information tracking the lead-up to impulsive choices is distributed across multiple frequency scales in STN, though current techniques may not possess sufficient signal-to-noise ratios to predict—and thus curb—impulsive behavior on a moment-to-moment basis.

Keywords: STN, DBS, impulsivity, balloon analog risk task (BART), decision making, local field potential (LFP)

INTRODUCTION

The well-attested success of deep brain stimulation (DBS) in ameliorating the symptoms of neurological disorders, such as Parkinson's Disease (PD) and essential tremor has brought with it an intense interest in DBS as a treatment modality for other conditions, including obsessive-compulsive disorder (OCD) and treatment-resistant depression (Johansen-Berg et al., 2008; Malone et al., 2009; de Koning et al., 2011; Figeo et al., 2013). This enthusiasm has also extended to an effort to identify biomarkers relevant to DBS targets for the cognitive symptoms of PD, in the hopes that alleviation of these symptoms

could also be addressed by next-generation DBS devices (Berney et al., 2002; Hälbig et al., 2009; Abosch et al., 2011; Hoang et al., 2017). These impulsive dimensions of the disease are thought to arise from the interaction between cortical planning areas, such as pre-SMA/SMA and the subthalamic nucleus (STN), the site of DBS stimulation, via the so-called “hyperdirect” pathway (Nambu et al., 2002; Aron and Poldrack, 2006; Frank, 2006; Cavanagh et al., 2011; Zavala et al., 2014). Studies have shown that oscillations in the theta band (4–8 Hz) play a key role in mediating the influence of frontal cortical areas on decisions in the basal ganglia and likely represent cognitive, top-down aspects of the decision process (Frank et al., 2007; Ballanger et al., 2009; Cavanagh et al., 2011; Zavala et al., 2014, 2015). This complements an extensive literature on the role of beta band oscillations (13–30 Hz) in the pathophysiology of Parkinson’s Disease, where it is known that such oscillations are suppressed by dopaminergic medication and effective DBS stimulation (Levy et al., 2002; Bronte-Stewart et al., 2009; Jenkinson and Brown, 2011; Whitmer et al., 2012). Similarly, power at these frequencies is reduced prior to voluntary movement, suggesting that such suppression is necessary for movement initiation (Jenkinson and Brown, 2011; Brittain et al., 2012).

However, while beta band activity has shown promising results as a biomarker suitable for training closed-loop DBS devices (Feng et al., 2007; Rosin et al., 2011; Santaniello et al., 2011), much less is known about the potential efficacy of stimulation targeting theta band. This stems, at least in part, from the more complex modulatory role posited for these oscillations in decision-making, where higher theta band power does not drive impulsivity but rather interacts with motivational signals entering the basal ganglia from the striatum and cortex (Cavanagh et al., 2011; Cavanagh and Frank, 2013; Zavala et al., 2014, 2015). Moreover, typical studies of these conflict-driven signals only report group averages, leaving open the question of variability across subjects and the robustness of these effects at the single subject level. Thus, the question of whether these signatures can be used and exploited at the level of individual subjects to predict and improve impulsive decisions in the face of conflict is a key one for the feasibility of next-generation DBS devices aiming to target and minimize cognitive side effects.

Here, we asked whether information about upcoming decisions might be disparately encoded by oscillations across the frequency spectrum by analyzing multi-channel intracranial recordings performed while patients with Parkinson’s Disease underwent implantation of DBS. We asked patients to perform a version of the balloon analog risk task (BART) (Lejuez et al., 2002), which is often used to model impulsive choice behavior (Lauriola et al., 2013). In contrast to standard perceptual decision tasks (cf. also Zavala et al., 2014), decisions in our experiment took place over a time scale of several seconds, allowing us to track neural correlates of the deliberative process as they occurred preceding decisions. We report that both single units and LFP signals altered their activity over time during the decision period, with LFP power in multiple frequency bands ramping over the course of the final seconds preceding decision, with phasic changes following the response. We also found evidence at the group level for separate roles for theta band activity in regulating

the decision process and beta band power in motor execution of the decision, consistent with prior studies. However, these patterns displayed wide variability at the group level, and a classifier trained to predict upcoming decision times failed to improve with the addition of field potential data. Together, these findings imply that decision-related information in the face of conflict is encoded at multiple frequencies within the basal ganglia, but development of additional, more reliable biomarkers may be needed as critical input for next-generation DBS devices designed to target impulsive decision-making.

MATERIALS AND METHODS

We conducted intraoperative recordings in 15 patients undergoing implantation of DBS devices in the subthalamic nucleus for treatment of Parkinson’s Disease. This study was carried out in accordance with the recommendations of the Duke University Medical Center Institutional Review Board with written informed consent from all subjects. All subjects gave written informed consent in accordance with the Declaration of Helsinki. The protocol was approved by the Duke University Medical Center Institutional Review Board. Patients were introduced to the study by a member of their treatment team. If interested, they were given further information by a study coordinator separate from the research team and afforded time to discuss with family prior to consent. Patients gave written consent preoperatively, either the morning of surgery or during a prior clinic visit. All patients who met the inclusion/exclusion criteria for DBS surgery were offered the opportunity to participate.

Surgical Procedures

All patients participating in the study ($N = 15$, 5 female, 10 male) underwent DBS implantation in the subthalamic nucleus (STN) for treatment of Parkinson’s Disease (see **Table 1** for details). Patients were off Parkinson’s medication (i.e., sinemet, dopamine agonists) for >12 h preceding surgery and during the surgical procedure to enhance their symptoms as part of the routine clinical care.

Patients underwent detailed MRI scanning, typically followed by Leksell frame placement, with a high resolution CT scan used to localize the surgical target. A subset of patients underwent frameless implantation following placement of fiducial markers in the skull later referenced to MRI scans. In all cases, STN was identified by first mapping out indirect target coordinates (i.e., $X = 11$ – 12 mm off the midline, $Y = 2$ mm behind AC-PC midpoint, $Z = 4$ mm below AC-PC line) and then refining these initial coordinates by the area of decreased FLAIR density on sagittal and coronal reconstructions corresponding to STN, with small adjustments (typically 1–2 mm) in X and Y dimensions for best fit. Single-unit recordings were first performed to define the borders of the STN, according to standard electrophysiological criteria, with the goal of attaining at least 5.5–6 mm of typical STN multi-unit firing. Typically, 2–3 passes were performed. Localization was performed using single-channel tungsten microelectrodes (Frederick Haer, Brunswick, ME).

TABLE 1 | Patient characteristics.

Age	Disease duration (years)	LEDD	Surgery side
66–70	2.5	245	Unilateral, R
61–65	10	2,930	Bilateral
56–60	3	700	Bilateral
61–65	7	1,098	Bilateral
66–70	22	800	Bilateral
61–65	12	898.5	Bilateral
66–70	5	925	Bilateral
56–60	6	560	Unilateral, R
46–50	5	0	Unilateral, L
51–55	6	965	Bilateral
41–45	5	0	Unilateral, L*
66–70	14	510	Bilateral
56–60	9.5	1674.5	Bilateral
51–55	13	991	Bilateral
61–65	4	800	Bilateral

*Previous R sided DBS.

M: male; F: female; LEDD: levodopa equivalent daily dose; R: right; L: left.

In most patients, data were collected during the process of STN localization. In some patients, data were collected from both sides of a bilateral DBS implantation (subjects 14, 16, 17). Once a well-isolated single unit was identified, the electrode would be left in place as the patient performed the behavioral task. Once data collection was complete, STN mapping would resume. In a subset of patients ($N = 7$; datasets 16.2, 17.2, 18.1, 20.1, 22.1, 23.1, 30.1), after full STN mapping, a 32-channel Pt/Ir microwire (35 μm diameter) array (Ad-Tech Medical Instruments, Racine, WI) was passed to the STN via an outer cannula (Patil et al., 2004) to a depth at which significant activity had been noted during initial localization. After allowing a few minutes for initial recordings to stabilize, the microwire array was slowly advanced through the cannula. Channels corresponding to putative single units were then selected and sorted, following which patients performed the behavioral task.

Following completion of the behavioral task, the DBS treatment electrodes were implanted. Patients undergoing bilateral DBS implantation were offered the opportunity to participate in the research twice, once for each implantation target. As a clinical routine, a brain computed tomography scan was performed within 12 h of the surgery procedure, and in no instance was a hemorrhage or other complication noted. Hence, the clinical risks of temporary placement of the 32 channel microwire array were demonstrated to be very low, as previously reported (Patil et al., 2004; Hanson et al., 2012).

Behavioral Task

We used a continuous time version of the BART (Lejuez et al., 2002), in which subjects must balance risk and reward in an attempt to maximize their score. In the original BART subjects choose whether to inflate a balloon, increasing its point value but also its risk of popping, or stop inflating and thereby add the current value of the balloon to their accumulated score. In

the original task, each balloon had a maximal number of times it could be inflated (e.g., 10 or 20) and an equal chance of popping on each trial. If the balloon popped, a subject would receive no points for that balloon.

Our variant modified this task to the constrained setting of the patient population in the operating theater. We designed a continuous-time version of the BART in which patients pressed a button to start inflating each balloon and pressed the button a second time to stop inflating and add its point value to their score (**Figure 1A**). In this version, balloons were divided into three risk levels, indicated by color. Each color corresponded to a distribution of maximal inflation (pop) times, with standard deviation proportional to the mean. This version proved much easier for patients to learn (pop times are more predictable) while preserving the cost/benefit tradeoff of the original task. Each balloon began with a value of 0 points, with its value increasing by 30 points/s. Pop times were drawn from a normal distribution on each trial (red: mean: 3 s, std: 0.9 s; orange: mean: 6.5 s, std: 1.95 s; yellow: mean: 10 s, std: 3 s). When patients successfully stopped the trial, the balloon turned green, a cash register sound played, and the balloon's point value was added to the accumulated score. If patients did not stop the trial in time, the balloon popped, resulting in a flash and a pop sound, and no points accrued.

For a subset of patients, two types of control conditions were added. In the first, 25% of balloons constituted forced trials. On these trials, a gray ring appeared at a fixed distance from the balloon, its radius drawn from the balloon's color-specific pop time distribution. That is, patients were explicitly shown a representative pop time of the balloon. On these trials, patients could not stop the balloon's inflation, and when the balloon's radius met the gray ring, the cash register sound played and patients received the current point value of the balloon. This subset of trials thus informed patients about the statistics of the task and attempted to control for motor anticipatory effects—since the outcome of the trial is certain and there is no need for stopping movement, motor planning should not play a role.

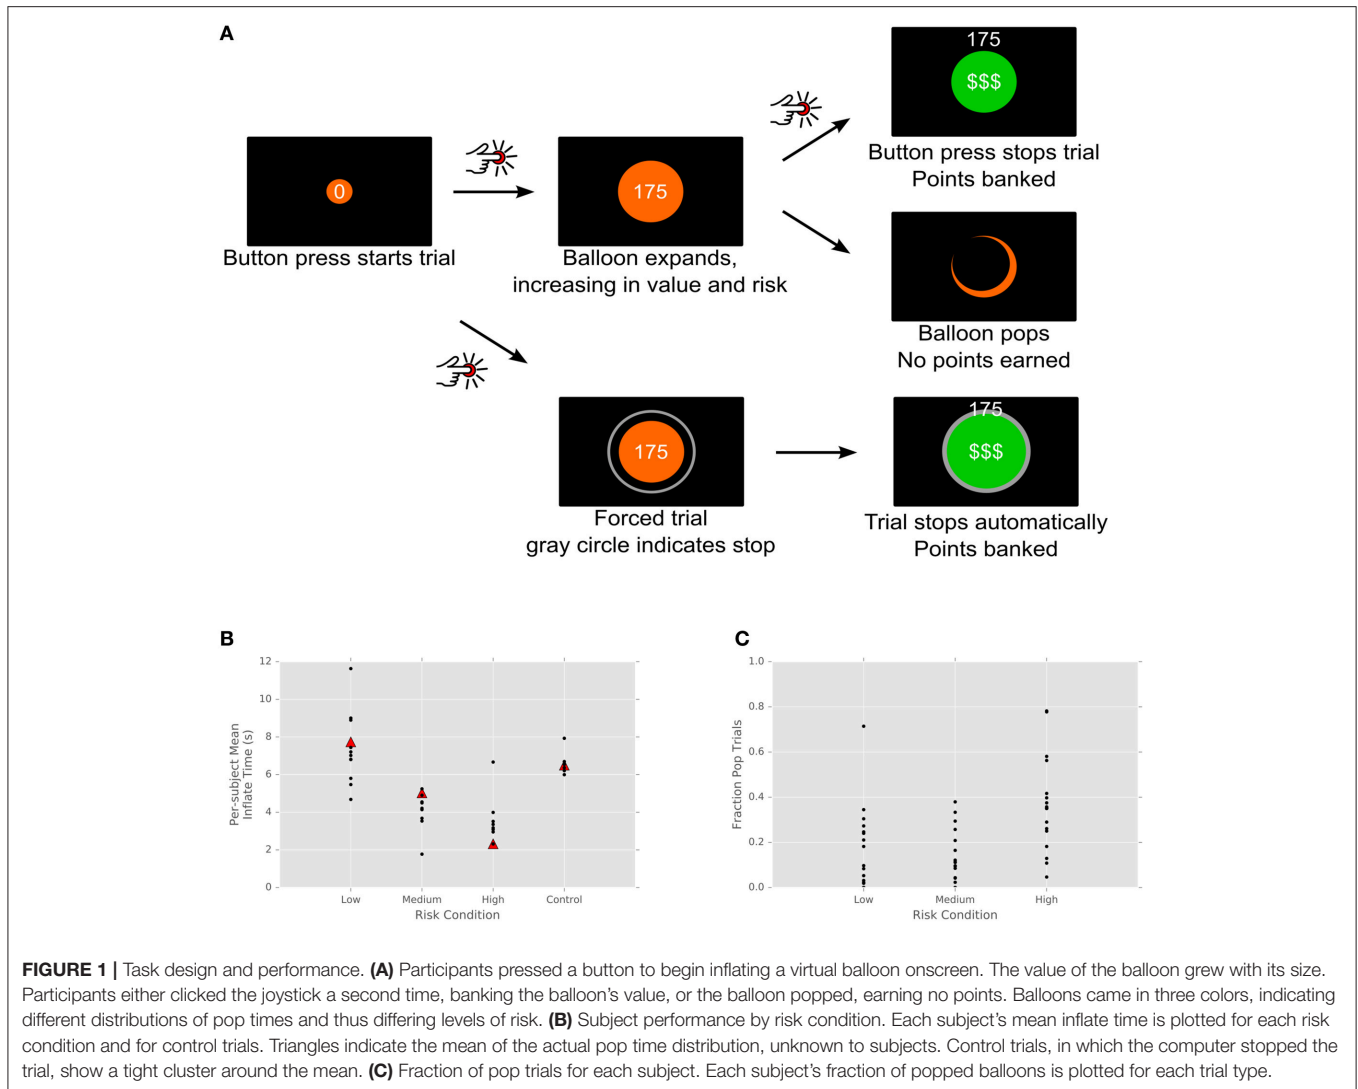
The second type of control trials involved gray (no-reward) balloons. Like the first type of control trial, these gray balloons were surrounded by a gray ring (with radius drawn from the intermediate/orange distribution). They functioned exactly as the first type of control trials, save that these balloons were worth no points. As a result, these trials control for effects of reward anticipation in the task.

Subjects were paid based on performance in the task, up to a maximum of \$40. Subjects were informed in advance of the rank order of risk for each balloon color (yellow, orange, red), to minimize the need for within-task learning. Subjects were also given the opportunity to familiarize themselves with the task prior to the experimental session.

The task itself was coded in Matlab (The Mathworks, Natick, MA) using PsychoPhysics Toolbox (Brainard, 1997). These scripts are available online at <https://github.com/jmxpearson/bart>.

Data Acquisition and Preprocessing

Data were recorded using a Plexon MAP system (Plexon, Inc., Dallas, TX) and the FHC Guideline 4000 (FHC, Inc.,



Bowdoin, ME). For single electrode recording, both high-pass and low-pass (local field potential; LFP) signal were recorded, with spikes extracted from the high-pass signal offline (see below). For multielectrode recording, LFP was recorded from all 32 channels and high-pass signal from the 16 most active channels, though all spike detection, and sorting was performed offline.

Prior to analysis, spike detection and sorting were performed via the WaveClus package (Quiroga et al., 2004) using a 4 standard deviation threshold based on robust (median-based) estimation of the noise. Putative units with <5 or >100 spikes/s average firing over the session were rejected based on established physiology of STN. Single units were identified based on reliability of spike waveforms (Tankus et al., 2009), with the remainder of units classified as multi-unit activity. Local field potentials (LFPs) were downsampled to 200 Hz. Recording artifacts were identified via either an aberrant increase in power or a sustained maximal amplitude (“railing”) and epochs containing these artifacts removed in subsequent analysis.

Statistical Analysis

Behavioral Analysis

To calculate optimal behavior in the task, we made use of the fact that pop times in each condition were Gaussian distributed with known mean and variance, while rewards were proportional to elapsed time. That is, $\tau \sim N(\mu, \sigma^2)$ and $R = \alpha t$, which implies an expected reward at time t proportional to both the elapsed time and the survival probability:

$E[R|t] = \alpha t(1 - CDF(t; \mu, \sigma^2))$ with $CDF(t)$ the cumulative distribution function of the Gaussian with given mean and variance. For each risk level, we then numerically calculated t_{opt} by maximizing expected reward.

Spiking Model

We modeled spike data as arising from a Generalized Linear Model with Poisson distribution and log link function:

$$N \sim \text{Poisson}(f \Delta t)$$

$$\log f = \beta_0 + \beta X + \gamma \log(t - t_0 + 1)$$

That is, counts in each $\Delta t = 50$ ms bin were assumed to be Poisson distributed with rate f . Task variables were modeled as contributing linearly to the logarithm of f , such that the baseline firing rate of the cell is $f_0 = e^{\beta_0}$ and the effect of adding a binary regressor x_i is a multiplicative gain e^{β_i} . As a result, our model predictors were time series, incorporating each bin for the entire experiment into a single regressor.

In addition to baseline spiking rate, our model includes an indicator term for the entire trial period (equivalent to the contrast between trial and inter-trial baseline), separate binary indicators for the inflation and outcome periods, a categorical regressor for outcome type (banked reward vs. no points), and a categorical regressor for trial type (risk levels high, medium, low, and control). Note that these regressors are not orthogonalized and suffer from multicollinearity, which we treat using regularization methods (see below). Finally, we include an effect of elapsed time within each trial, such that if t is the elapsed time and t_0 the trial start time, the firing rate increases as $f = \alpha(t - t_0 + 1)^\gamma$, with α including the effects of all other variables.

We fit this model using elastic net regularization via the `glmnet` package in R (Friedman et al., 2010). In our results, we used 10-fold cross-validation to choose the values of λ and α . That is, we fit the model using 90% of the available data (time bins) and used the remaining 10% of data to assess its predictive performance. We then repeated this method 10 times using distinct random partitions of the data. For λ , we used the largest value (returned by `glmnet`) within one standard deviation of the value that produced the best performance, on average, for the cross-validation procedure. That is, we used the most parsimonious model (highest λ) statistically indistinguishable from the best-performing model. To determine α , we repeated the entire `glmnet` fitting 11 times, testing values of between 0.1 and 1 in steps of 0.1.

In addition, for comparison, we also calculated effect sizes for other regularization choices. These include the same model without regularization (more exactly, set to the minimum value used by the elastic net algorithm), with the value of λ producing the best predictive performance, and the parsimonious model described above fit to shuffled data. In the last case, spikes counts for each time bin were randomly shuffled, removing any temporal correlations in the data that might be explained by our regressors.

Local Field Potentials

LFPs were preprocessed by first subtracting the mean across all channels at each time point from each channel to remove global artifacts, followed by bandpass filtering into typical frequency bands of interest as follows (frequencies in Hz): delta: (0.1, 4), theta: (4, 8), alpha: (8,13), beta: (13, 30), gamma: (30, 80). These signals were subsequently downsampled to 200 Hz and instantaneous power calculated via the square of the Hilbert transformed signal. Finally, time series for each channel were censored to remove microphonics, amplifier saturation, and other artifacts.

For peri-event LFP analyses, summary traces were computed by collecting samples around each event of interest (trial

starts, trial stops, etc.), taking the median at each time point across events, and (for plotting purposes) standardizing the log power for each channel to have mean zero and unit variance. This last step allowed us to compare dynamics of the signal across channels despite variability in channels' dynamic range. Broadband signals were processed in the same manner, omitting the step of bandpass filtering. In some plots, traces were subsequently averaged across channels on the logged and z-scored scale. For instantaneous power analyses that collapsed across subjects, the group mean of subject means on the logged and z-scored scale was used. By normalizing prior to averaging, both of these choices give equal weight to each channel, rather than favoring channels with higher overall power. However, we also consider the case of average across channels without z-scoring in Supplementary Results.

Time-frequency plots were constructed via continuous wavelet transforms of the data. For analyses that collapsed across channels and subjects, individual channels were z-scored prior to averaging, so that channels with widely varying power contributed equally. We used a complex Morlet wavelet with $w = 9$ and logarithmically-spaced frequencies from 2.5 to 50 Hz. Power was calculated via the sum of squares of the real and imaginary parts of the wavelet transform and plotted on a decibel scale. To compensate for different levels of intrinsic power within frequency bands, the plotted signal was normalized by frequency band by the mean power present across all trials within that band in the interval $B = (-1.5, -0.5)$, that is, from 1.5 s before to 0.5 s before the event of interest. Specifically, the plotted intensities are given by $I = 10 \log_{10} \frac{P(f,t)}{P_0(f)}$ with $P(f, t) = |s(f, t)|^2$, $s(f, t)$ the wavelet-transformed signal, and $P_0(f)$ the mean power in the baseline interval. For time-frequency plots depicting the contrast between two conditions, no normalization was performed, so that the plotted intensities reflect $I = 10 \log_{10} \frac{P_1(f,t)}{P_2(f,t)}$, the ratio of power between the two conditions in decibels.

For contrast time-frequency plots, we determined significant areas of relative activation via a cluster-based bootstrap algorithm (Nichols and Holmes, 2001). For each pixel, we calculated an F-statistic between the (decibel-scaled) power for the two conditions, using a 5×5 pixel local smoother to improve estimation of variance. We used p -value thresholds of 0.1 and 0.9 for these F-statistics to determine local clusters of activation. That is, pixels in the uppermost and lowermost deciles of cumulative probability density were selected for further analysis. Groups of contiguous pixels exceeding this threshold constituted a cluster, with cluster mass equal to the sum of log F-values for all pixels in the cluster. Distributions of this mass statistic under the null hypothesis were generated by randomly shuffling the labels of the two conditions and then repeating the entire analysis pipeline above (smoothing, F-statistic, thresholding, clustering). All clusters with mass statistic significant at the $p < 0.05$ threshold with respect to this null distribution were retained.

The LFP prediction algorithm made use of a modified sparse proportional hazards model. Inputs to the classifier were determined as follows: for each dataset, LFP in each channel was bandpass filtered into delta, theta, alpha, beta, and gamma bands as defined above, then decimated to 40 Hz, converted

to instantaneous power via the Hilbert transform, censored for artifacts, log scaled, and standardized to 0 mean and unit variance. Regressors were the logged mean power for each (channel, band) combination for 0.5 s preceding each time point of interest. True positive events (coded as 1 s) were chosen as the response time for trials in which participants successfully stopped balloon inflation and received a reward. All other data points were non-events (coded as 0 s). We assumed these binary counts to be Poisson distributed within each time bin with instantaneous rate (at risk level i):

$$\lambda_i(t) = h_i(t) \exp(\mu + X(t) \cdot \beta)$$

with $h_i(t)$ the hazard function for the normal distribution with (fitted) parameters m_i and s_i , $X(t)$ the regressors (LFP power in each band, channel) at time t , and μ and β model parameters distributed according to a horseshoe prior (Carvalho et al., 2009). The utility of the horseshoe prior lies in its privileging of sparse solutions for the regression weights, allowing us to simultaneously perform inference and model selection. That is, channels and frequency bands are parsimoniously shrunk to 0, as in the LASSO model above. Note in particular that the use of instantaneous log power for $X(t)$ implies that the event rate is the equal to the hazard rate modulated by a weighted geometric mean of momentary power.

We fit these models using the Stan probabilistic programming language (Carpenter et al., 2016) using four chains of 2,000 iterations each with a burn-in period of 1,000 samples (half the total) and a thinning fraction of five. Convergence was assessed by means of effective sample sizes and the Gelman-Rubin statistic (Gelman and Rubin, 1992).

RESULTS

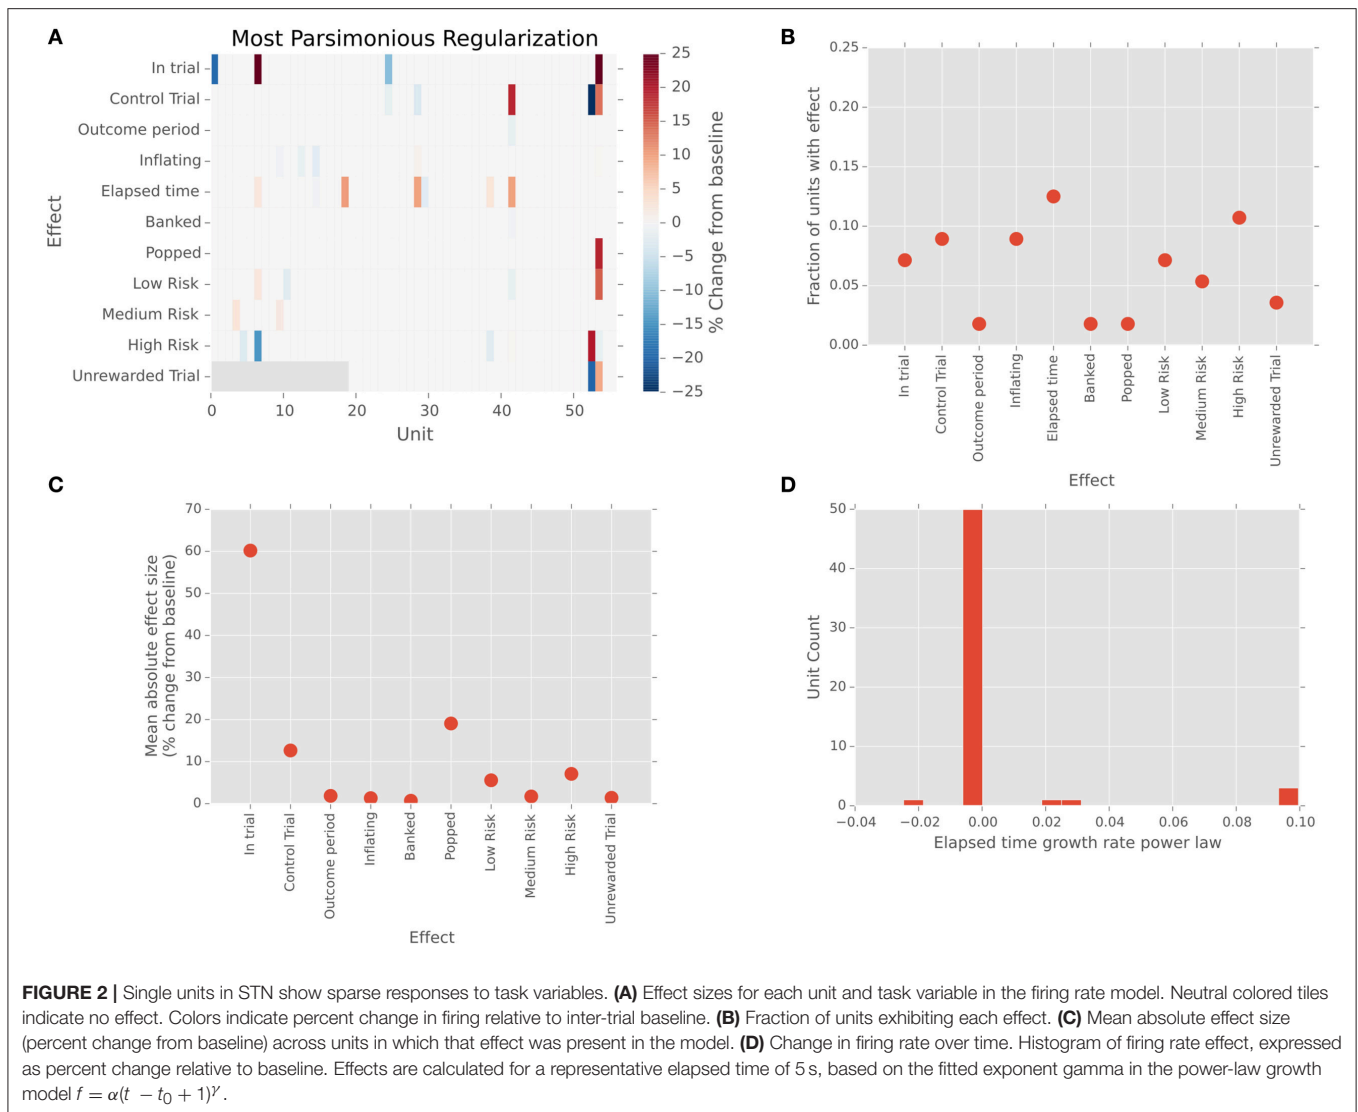
As expected, subjects balanced risk and reward, with longer inflation times for less risky balloons. In addition, nearly all subjects exhibited risk aversion across each of the categories, with mean trial times below optimal for the low and medium risk conditions (but above in the high risk condition; **Figure 1B**). This is consistent with behavior observed in the standard BART (Lejuez et al., 2002; Lauriola et al., 2013), where the level of risk aversion is the measure of interest, since it has been shown to correlate inversely with real-world measures of risk-taking. This is also consistent with **Figure 1C**, which shows the fraction of trials for each subject that resulted in no points (i.e., popped balloons). In the standard BART, where the total number of trials is fixed, optimal strategy dictates that subjects should inflate balloons to half their maximal size—that is, on average, half of trials should result in pops. In our modified paradigm, by contrast, the number of trials was time-limited, with the result that subjects were incentivized to adopt a strategy of shorter, more numerous trials with higher probability of reward. However, what is most important for our subsequent analysis of neural data is not whether subjects pursued behaviorally optimal strategies but the fact that varying risk levels in the task resulted in a range of decision times, varying (roughly) between 1 and 12 s.

Most importantly, these responses were self-initiated. Rather than rapid responses to external cues, they were the result of deliberative behavior, permitting us to examine the time evolution of signals underlying the process of decision formation within STN.

We recorded 56 ($N = 48$ single, 8 multi-) units from our patient population. All but one of these units was determined to lie within STN boundaries based on results of the intraoperative mapping process and pre-surgical MRI-based targeting. (One unit, with regular baseline firing rate of 150 spikes/s and an absence of bursting, was most likely sampled from the nearby substantia nigra pars reticulata.) We analyzed spiking activity of the STN units via a regularized, cross-validated GLM approach that incorporated all task variables into a single model. **Figure 2A** shows the effect sizes for each unit and task variable, color-coded as a percent change from baseline firing. Individual cells show differences in firing between risk levels, as well as between reward and no-reward trials and between standard (free choice) and control (no reward) trials. That is, local firing rates are sensitive to risk and the reward, as well as to the difference between trials on which a motor action must be made and those in which passive viewing alone is required. Clearly, the effects are sparse throughout the population, as is typical for neurons in basal ganglia. However, it is also important to note that the regularization and cross-validation approach we have used is extremely conservative; a similar approach using the best-performing (not most parsimonious) regularization found similar effects much more strongly represented, with roughly 80% of cells exhibiting each effect. Along the same lines, a shuffle-corrected version of the same analysis produced zero false positive responses (Figure S1). That is, multiple checks indicate that the sparse effect matrix above represents robust effects in our data, not statistical artifacts of our analysis.

We found neurons responsive to a variety of task-related variables, including those active for the difference between task and baseline ($N = 4/56$, 7%) and in the contrast between risk levels (**Figure 2B**; $N = 11/56$, 20%). We likewise found a range of effect strengths, with the effect of task vs. baseline and popped balloons largest (**Figure 2C**). Most intriguingly, the most commonly represented effect in our neural sample was the effect of elapsed time on neural firing. Roughly 14% ($N = 8/56$) of cells in our population showed small to moderate changes in firing rate over the course of balloon inflation, with effects on the order of a few percent to 20% changes in gain for a representative inflation time of 5 s (**Figure 2D**). These findings are consistent with previous reports of increased oscillatory power in STN in tasks requiring rapid responses (Cavanagh et al., 2011; Frank et al., 2015), but to our knowledge, this is the first single unit evidence of this mechanism operating over longer decision timescales in humans.

Intrigued by these findings in single units, we looked for further evidence of ramping activity at the population level preceding decisions. LFP are thought to measure voltages due to the summed activity of neurons in a local region around the electrode tip (potentially including some long-range effects) (Kajikawa and Schroeder, 2011). As such, they offer a glimpse into the rate and magnitude of local neural processing in a

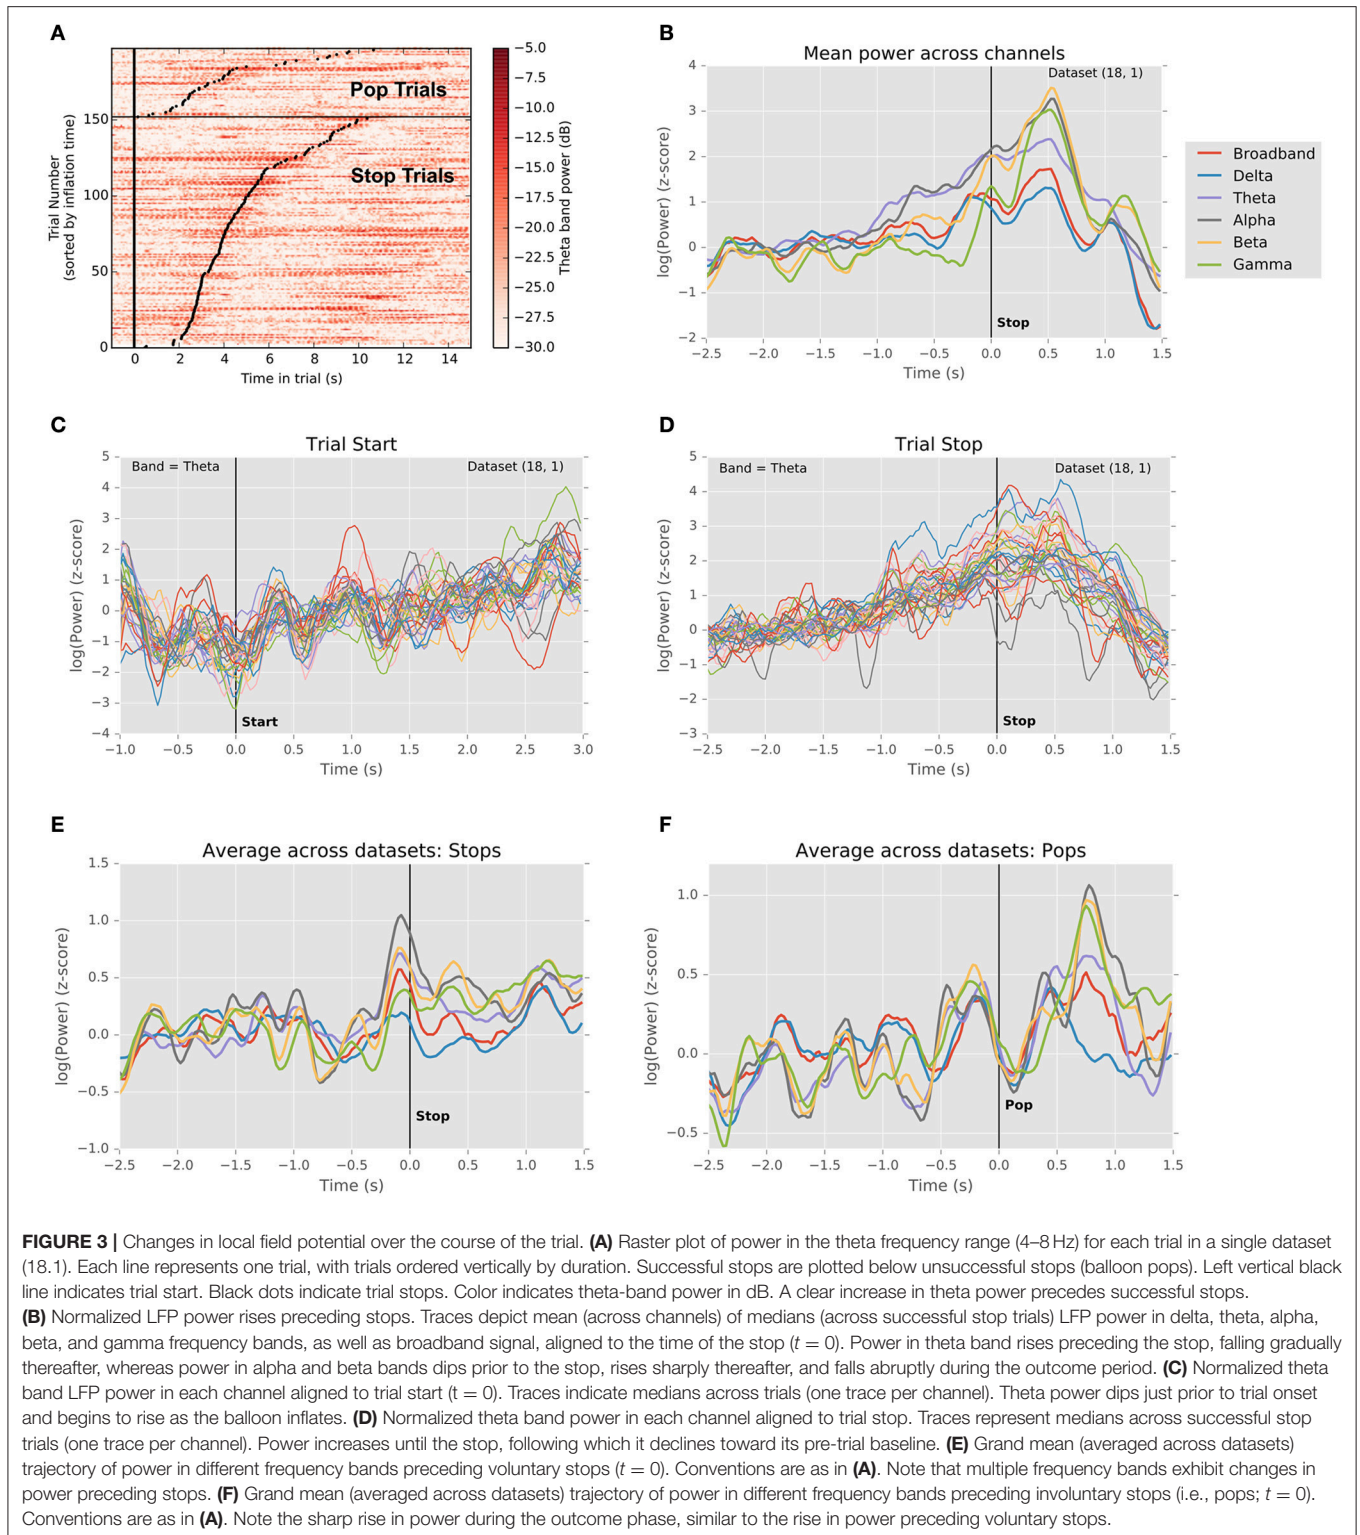


region. We recorded LFPs simultaneously with our single units, in some cases from up to 32 channels simultaneously. We asked whether these signals contained information about upcoming stop decisions, and whether patterns of local activity might serve to predict these decisions in a data-driven way.

Figure 3 explores this question by depicting LFP power in the theta frequency band (4–8 Hz) as a function of time on each trial for a single subject. Theta band activity measured by scalp EEG has been shown to be correlated with reaction times in a learned stimulus association task (Cavanagh et al., 2011; Zavala et al., 2013, 2014), where it was hypothesized that this is due to the influence of medial frontal cortex on integration of information in STN via the hyperdirect pathway. Here, in a self-paced decision context, we see a similar pattern recapitulated in theta-band power within the STN. Each line of the image in **Figure 3A** represents a single trial. The vertical line marks the start of each trial, while each black dot marks its end. Trials have been grouped by outcome, and ordered by duration. For successful stops, theta band power clearly increases preceding stop movements, with

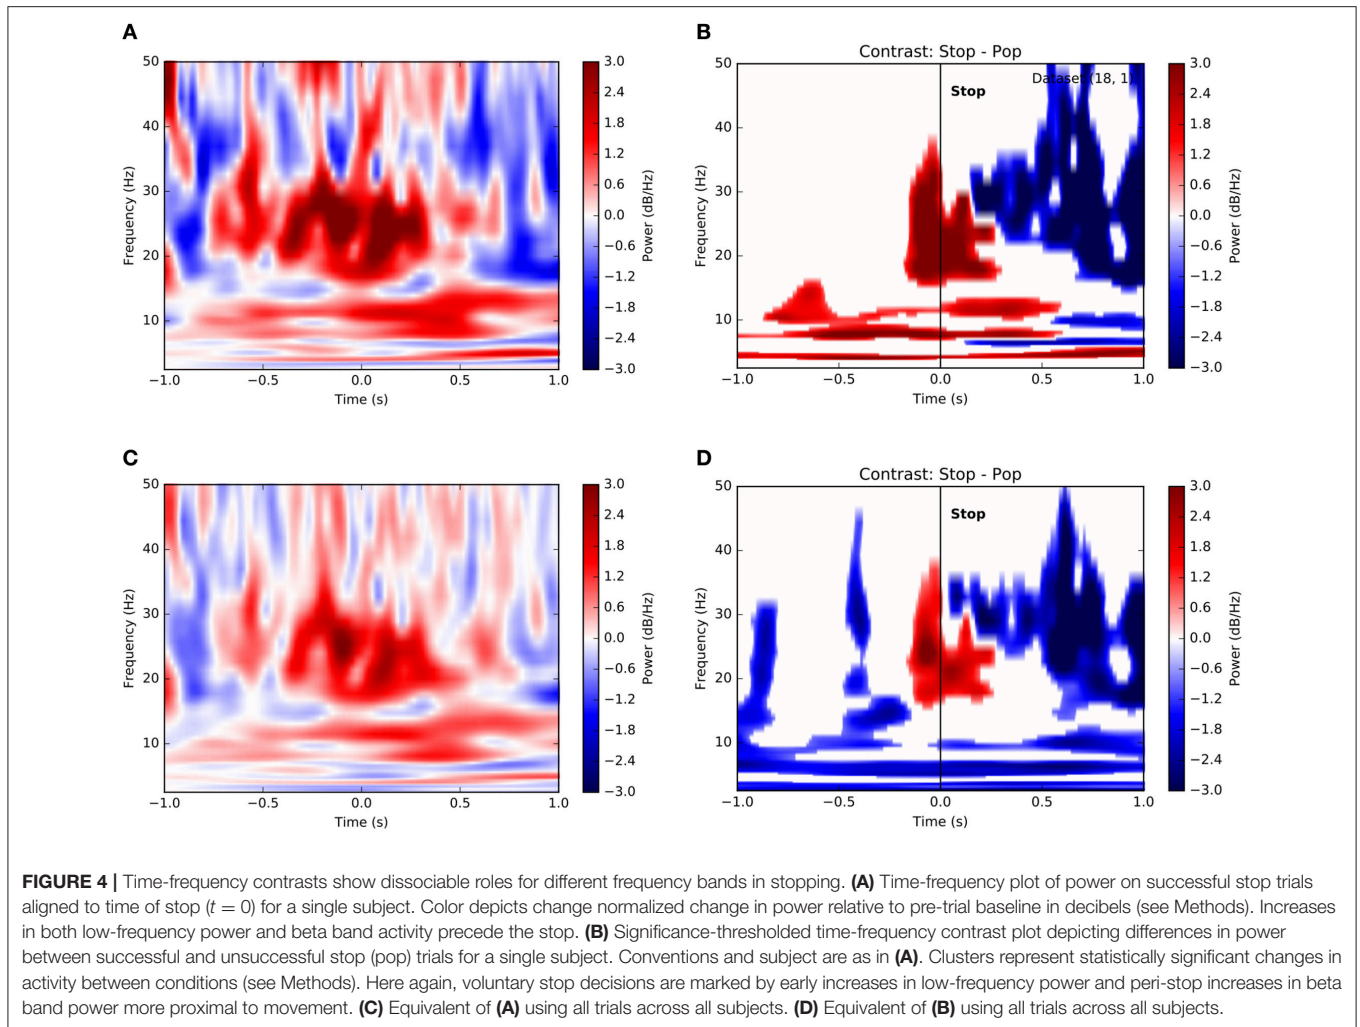
some activation lasting into the outcome period. By contrast, this pattern is largely absent prior to unexpected stops (popped balloons). Like the data from single units, this suggests that local activity in STN begins to rise during the decision interval, peaking around the time of movement.

Figure 3B extends this analysis to other frequency bands, depicting the trial- and channel-averaged power surrounding trial starts and successful stops for the same subject as in **Figure 3A**. Clearly, power ramps over the 2 s preceding stops across frequency bands ($N = 7/7$ subjects, $p < 0.05$, t -test for regression slope), with alpha and theta power exhibiting a sustained rise. Intriguingly, for this subject, beta and gamma power begin their rise relatively late, <500 ms prior to the stop and coincident with a local peak in delta power. This is consistent with the idea that, for instance, beta band power increases precede motor movement (Jenkinson and Brown, 2011; Brittain et al., 2012) and that theta band power is related to a threshold for evidence accumulation (Cavanagh et al., 2011; Frank et al., 2015), though clearly peri-event activity is not confined to only



a single frequency band. In **Figures 3C,D**, for the specific case of theta band, we see that this rise in power begins well before, at trial onset, and that a dip in power precedes the initiation of the trial. As a result, rises in theta begin well before the time of movement, suggesting that this process pertains to planning, not

simply movement initiation. Likewise, at the population level, ramping activity preceded stops, and sharp increases in power followed pops, consistent with previous findings (**Figures 3E,F**; Cavanagh et al., 2011; Tan et al., 2014; Zavala et al., 2014, 2015).

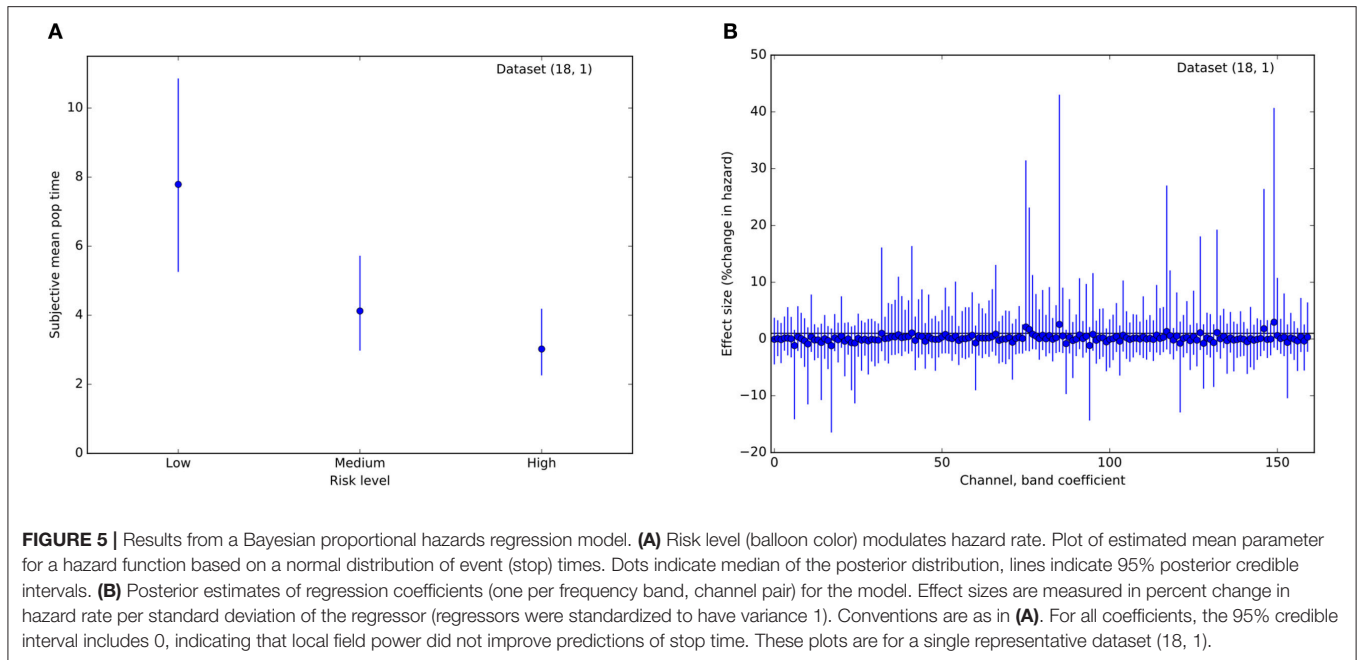


Observing that population averages in **Figures 3E,F** exhibited ramping activity across multiple frequencies, we asked whether such an integrative signal corresponded specifically to theta-band activity or whether it might involve oscillations at multiple scales. **Figure 4A** is a time-frequency plot showing channel- and trial-averaged LFP power preceding successful stops for a single subject. Clearly, power increases both at frequencies below 15 Hz and within the beta band in the second preceding decision, peaking just after decision at the time of reward. Indeed, this pattern remains when we consider the contrast between successful stops and unanticipated stops (popped balloons) (**Figure 4B**). Again, significant increases in low-frequency (<15 Hz) and beta band power distinguish between the two conditions, suggesting that decisions to stop require an increase in theta- and alpha-band activity along with a concomitant increase in beta power, as has been reported elsewhere (Bronte-Stewart et al., 2009; Jenkinson and Brown, 2011; Whitmer et al., 2012; Bastin et al., 2014). In fact, very similar results hold for a population-level analysis based on the average across all trials and all subjects (**Figures 4C,D**).

It may be considered surprising that **Figure 4** depicts an increase in beta-band power, as opposed to the decrease reported

in, e.g., Cavanagh et al. (2011) and elsewhere. In fact, this results from our choice of normalization. In **Figure 4**, we have first normalized the variance within each recording channel prior to averaging, which treats each signal on an equal footing. However, a raw average across, without normalization, shows a decrease in beta around the time of stops (Figure S3). This suggests that the largest modulations of LFP power in STN are negative, though weaker positive modulations may be widespread.

There is also suggestive evidence for this interpretation from a subset of participants who performed control conditions with both voluntary and forced stops. Figure S2A shows that in this case, voluntary stops, which involve both decision and motor movement on the part of participants, exhibit greater power in the beta and theta bands preceding stops. That is, the LFP effects noted above are related to the decision and movement process, not simply the temporal structure of the trial. Similarly, when comparing two identical motor actions, the movements at trial start and trial stop (Figure S2B) the differences in power between the two conditions appear to be quite similar to the overall activity for end-of-trial stops, suggesting that these effects pertain to the decision process itself. Thus, converging evidence supports the notion that gradual increases in low-frequency power are



reflective of the deliberative decision process, while beta power changes precede movement initiation, followed by a rapid rise in beta post-movement and a fall across all bands after reward.

Therefore, we asked whether activity across multiple channels might collectively be harnessed to predict trial stops at the single-subject level. Using a survival analysis approach, we modeled the time to event (voluntary stop) as dependent upon an increasing hazard rate. Trials on which the balloon popped were treated as censored data, with the voluntary stop time unobserved. Most importantly, we assumed that the instantaneous probability of stopping was modulated by a weighted geometric mean of momentary power within each frequency band and channel. That is, the probability of a stop increased with time in a manner specific to each balloon color, modulated by a weighted combination of local neural activity. As depicted in **Figure 5A**, models for most subjects exhibited different hazard rates for the different risk levels, resulting in earlier stops for riskier balloons. However, with only a single exception, the coefficients encoding neural effects on stopping were not appreciably different from 0 for any subject, frequency band, or channel (95% credible interval for β overlapping 0; e.g., **Figure 5B**). Thus, while STN activity appears to track the deliberative process by ramping power across multiple frequency bands, this activity does not appear to be specifically predictive of the stop decisions themselves on a single-trial basis.

DISCUSSION

By recording directly from humans undergoing invasive brain surgery, we have demonstrated not only the presence of single units within the basal ganglia that increase firing over the course of the decision process, but that population measures like the local field potential also exhibit ramping activity during the

process of deliberation. This finding is consistent with previous experimental (Frank et al., 2007, 2015; Cavanagh et al., 2011; Brittain et al., 2012; Zavala et al., 2013, 2014) and modeling (Frank, 2006) work positing a role for the mediofrontal-hyperdirect pathway in stopping behavior, but extends these findings beyond the realm of cued reaction time tasks to the longer timescale of deliberative decisions, confirming the applicability of experimental and modeling predictions to human cognitive tasks.

We likewise found single units and LFP signals responsive at the time of reward or negative outcomes, with stronger effects for balloon pops than rewards. This is consistent with single unit studies in animals showing responses to both rewards and error trials (Isoda and Hikosaka, 2008; Lardeux et al., 2009; Baunez and Lardeux, 2011), and likely reflects widespread neural activity in response to reinforcing outcomes across the brain. However, the most prevalent single unit response was to elapsed time in the trial, suggesting that STN not only plays a role in stopping behavior, but also encodes critical information related to upcoming choice during more prolonged deliberative processes.

Likewise, much debate has focused on the role and relative importance of oscillatory activity within STN and the basal ganglia more broadly. Activity within the beta band is known to be related both to upcoming motor movement, with desynchronization preceding movement onset (Jenkinson and Brown, 2011), and to the particular pathophysiology of Parkinson's Disease, where beta activity is known to be enhanced, while medications and effective DBS can suppress beta activity with symptomatic treatment (Bronte-Stewart et al., 2009; Jenkinson and Brown, 2011; Whitmer et al., 2012). In addition, theta band activity recorded in scalp EEG, thought to originate over medial frontal cortex, perhaps in pre-SMA/SMA, has been linked to cortical control over stopping in the basal

ganglia (Aron and Poldrack, 2006; Cavanagh et al., 2011; Zavala et al., 2013, 2014), consistent with single-unit studies in non-human primates (Isoda and Hikosaka, 2007, 2008). Previous work in a cued response task showed both theta enhancement and beta suppression prior to responses, with theta linked via computational modeling to decision threshold in a drift diffusion model (Cavanagh et al., 2011; Zavala et al., 2014; Frank et al., 2015). Our findings complement this work by demonstrating a gradual increase in LFP power across multiple frequency bands preceding self-paced decisions, reflecting widespread STN activity. Thus, while beta desynchronization does occur before motor movement, antecedents of the decision output occur within multiple frequency bands after the decision process is activated. Furthermore, we provide additional evidence for a posited dissociation between beta/motor and theta/cognitive inputs to STN, consistent with the hypothesis that these signals reflect decision-related information from cortical afferents, likely in primary, associative motor, and midline limbic areas.

Several previous studies have delineated a specific role for theta-band activity in STN as setting a threshold for stopping behavior in the context of the drift-diffusion model (Frank, 2006; Cavanagh et al., 2011; Zavala et al., 2013). Evidence suggest that this activity represents top-down inputs from medial prefrontal cortex, with stronger theta activity present in the case of increased conflict. However, in our data, theta activity clearly increases over the course of the trial, favoring a role for theta in tracking the decision process rather than setting a threshold. Several observations may explain this discrepancy. First, theta may indeed play a dual role in both the accumulation process and threshold-setting, one which our design did not have sufficient power to detect. Second, the distinct demands of our task, which required neither a speeded response nor a forced choice, might have resulted in qualitatively different patterns of network activity. Likewise, our task offered only a single response option, asking participants a question of “when” rather than “which,” and so likely failed to recruit mechanisms responsible for mediating response conflict. Finally, another study of response conflict in STN on somewhat longer timescales (Zavala et al., 2014) found no correlation between theta band activity and response time, consistent with our results. However, we are agnostic as to whether these early onset signals represent conflict or some other psychological construct, since many similar ideas like dread, risk, or reward anticipation may plausibly rise during the course of our task, and none constitutes a mechanistic explanation. Future studies will no doubt be needed to elucidate the specific function or functions of theta-band activity in STN.

However, there are several limits to the current study. First, as noted above, there remains some uncertainty as to the exact role played by frontal theta in conjunction with the STN. While the focus has often been on conflict and stopping, Brittain and Brown (2014) have also suggested a relationship with attention. Huebl et al. (2011) found a link between alpha band activity and emotional state. Indeed, the signals we have identified as part of the deliberative phase span frequency bands and are no doubt involved in multiple processes, many of which are difficult, if not impossible to dissociate in a single study. Finally, the positions of our microwire electrodes are distributed at random

throughout STN but may not sufficiently span the more medial areas associated with limbic and motivational inputs that we expect to contribute to impulsive decisions. Nonetheless, what is important for our purposes is that our task successfully elicits activity typically observed in studies designed to study impulsive decisions and so we are in a position to ask whether such signals are capable of predicting behavior at a single-trial level.

Most importantly, our results underscore the need for identifying additional high-quality biomarkers for closed-loop DBS systems aiming to treat impulsive side effects of PD. A number of closed-loop systems have focused on suppression of beta-band activity as a marker for motor pathophysiology, but an endogenous electrophysiological surrogate marker for impulsivity remains elusive (Jenkinson and Brown, 2011; Rosin et al., 2011). Even with a much higher channel count than possible with typical DBS devices, we were unable to effectively predict stopping (that is, to improve our probabilistic prediction of the stop time) on a single trial basis by incorporating neural signals from all channels and all frequency bands. Given the multiply attested relationship between beta band activity and stopping, this is perhaps surprising, but can perhaps be explained by two observations: First, our experiment used high impedance microelectrode wires as opposed to larger DBS contacts, resulting in less averaging and more spatially local signals. It is possible that this coarse-graining may be helpful in mitigating noise and deriving more predictive biomarkers in other DBS studies. Second, most studies report signals averaged over a larger number of trials. While reliable, such signals may be dwarfed by noise on an individual trial basis and thus of little utility for prediction. This single trial limitation could be overcome by averaging across several trials for improved predictions, which may still be feasible to implement in a slower cognitive feedback system. Finally, more sophisticated algorithms, equipped with larger training data sets like those being collected from already implanted DBS devices, may be able to extract more reliable signals from numerous cross-channel correlations. Clearly, further work is needed to identify such signals for the next generation of closed-loop DBS devices.

AUTHOR CONTRIBUTIONS

JP, PH, SL, MP, and DT designed research. JP, PH, SL, and DT collected data. JP analyzed data. JP, PH, SL, MP, and DT wrote the manuscript.

ACKNOWLEDGMENTS

The authors would like to thank the patients of Duke University Medical Center who generously donated their time to make this work possible. This work was supported by NIH grants R21NS066115 (DT), R21NS084176 (JP), and K01ES025442 (JP).

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2017.00646/full#supplementary-material>

REFERENCES

- Abosch, A., Gupte, A., Eberly, L. E., Tuite, P. J., Nance, M., and Grant, J. E. (2011). Impulsive behavior and associated clinical variables in Parkinson's disease. *Psychosomatics* 52, 41–47. doi: 10.1016/j.psych.2010.11.013
- Aron, A. R., and Poldrack, R. A. (2006). Cortical and subcortical contributions to Stop signal response inhibition: role of the subthalamic nucleus. *J. Neurosci.* 26, 2424–2433. doi: 10.1523/JNEUROSCI.4682-05.2006
- Ballanger, B., van Eimeren, T., Moro, E., Lozano, A. M., Hamani, C., Boulinguez, P., et al. (2009). Stimulation of the subthalamic nucleus and impulsivity: release your horses. *Ann. Neurol.* 66, 817–824. doi: 10.1002/ana.21795
- Bastin, J., Polosan, M., Benis, D., Goetz, L., Bhattarjee, M., Piallat, B., et al. (2014). Inhibitory control and error monitoring by human subthalamic neurons. *Transl. Psychiatry* 4:e439. doi: 10.1038/tp.2014.73
- Baunez, C., and Lardeux, S. (2011). Frontal cortex-like functions of the subthalamic nucleus. *Front. Syst. Neurosci.* 5:83. doi: 10.3389/fnsys.2011.00083
- Berney, A., Vingerhoets, F., Perrin, A., Guex, P., Villemure, J.-G., Burkhard, P. R., et al. (2002). Effect on mood of subthalamic DBS for Parkinson's disease: a consecutive series of 24 patients. *Neurology* 59, 1427–1429. doi: 10.1212/01.WNL.0000032756.14298.18
- Brainard, D. H. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436. doi: 10.1163/156856897X00357
- Brittain, J.-S., and Brown, P. (2014). Oscillations and the basal ganglia: motor control and beyond. *Neuroimage* 85(Pt 2), 637–647. doi: 10.1016/j.neuroimage.2013.05.084
- Brittain, J.-S., Watkins, K. E., Joundi, R. A., Ray, N. J., Holland, P., Green, A. L., et al. (2012). A role for the subthalamic nucleus in response inhibition during conflict. *J. Neurosci.* 32, 13396–13401. doi: 10.1523/JNEUROSCI.2259-12.2012
- Bronte-Stewart, H., Barberini, C., Koop, M. M., Hill, B. C., Henderson, J. M., and Wingeier, B. (2009). The STN beta-band profile in Parkinson's disease is stationary and shows prolonged attenuation after deep brain stimulation. *Exp. Neurol.* 215, 20–28. doi: 10.1016/j.expneurol.2008.09.008
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., et al. (2016). Stan: A probabilistic programming language. *J. Stat. Softw.* 20, 1–37. doi: 10.18637/jss.v076.i01
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2009). *Handling Sparsity via the Horseshoe*. AISTATS. Available online at: <http://www.jmlr.org/proceedings/papers/v5/carvalho09a/carvalho09a.pdf>
- Cavanagh, J. F., and Frank, M. J. (2013). Stop! Stay tuned for more information. *Exp. Neurol.* 247, 289–291. doi: 10.1016/j.expneurol.2013.01.023
- Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., et al. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nat. Neurosci.* 14, 1462–1467. doi: 10.1038/nn.2925
- de Koning, P. P., Figeé, M., van den Munckhof, P., Schuurman, P. R., and Denys, D. (2011). Current status of deep brain stimulation for obsessive-compulsive disorder: a clinical review of different targets. *Curr. Psychiatry Rep.* 13, 274–282. doi: 10.1007/s11920-011-0200-8
- Feng, X.-J., Greenwald, B., Rabitz, H., Shea-Brown, E., and Kosut, R. (2007). Toward closed-loop optimization of deep brain stimulation for Parkinson's disease: concepts and lessons from a computational model. *J. Neural Eng.* 4, L14–L21. doi: 10.1088/1741-2560/4/2/L03
- Figeé, M., Luijckes, J., Smolders, R., Valencia-Alfonso, C.-E., van Wingen, G., de Kwaasteniet, B., et al. (2013). Deep brain stimulation restores frontostriatal network activity in obsessive-compulsive disorder. *Nat. Neurosci.* 16, 386–387. doi: 10.1038/nn.3344
- Frank, M. J. (2006). Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Netw.* 19, 1120–1136. doi: 10.1016/j.neunet.2006.03.006
- Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., et al. (2015). fMRI and EEG Predictors of Dynamic Decision Parameters during Human Reinforcement Learning. *J. Neurosci.* 35, 485–494. doi: 10.1523/JNEUROSCI.2036-14.2015
- Frank, M. J., Samanta, J., Moustafa, A. A., and Sherman, S. J. (2007). Hold your horses: impulsivity, deep brain stimulation, and medication in parkinsonism. *Science* 318, 1309–1312. doi: 10.1126/science.1146157
- Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* 33, 1–22. doi: 10.18637/jss.v033.i01
- Gelman, A., and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Stat. Sci.* 7, 457–472. doi: 10.1214/ss/1177011136
- Hällig, T. D., Tse, W., Frisina, P. G., Baker, B. R., Hollander, E., Shapiro, H., et al. (2009). Subthalamic deep brain stimulation and impulse control in Parkinson's disease. *Eur. J. Neurol.* 16, 493–497. doi: 10.1111/j.1468-1331.2008.02509.x
- Hanson, T. L., Fuller, A. M., Lebedev, M. A., Turner, D. A., and Nicolelis, M. A. L. (2012). Subcortical neuronal ensembles: an analysis of motor task association, tremor, oscillations, and synchrony in human patients. *J. Neurosci.* 32, 8620–8632. doi: 10.1523/JNEUROSCI.0750-12.2012
- Hoang, K., Cassar, I. R., Grill, W. M., and Turner, D. A. (2017). Biomarkers and stimulation algorithms for adaptive brain stimulation. *Front. Neuro.* 11:564. doi: 10.3389/fnins.2017.00564
- Huebl, J., Schoenecker, T., Siegert, S., Brücke, C., Schneider, G.-H., Kupsch, A., et al. (2011). Modulation of subthalamic alpha activity to emotional stimuli correlates with depressive symptoms in Parkinson's disease. *Mov. Disord.* 26, 477–483. doi: 10.1002/mds.23515
- Isoda, M., and Hikosaka, O. (2007). Switching from automatic to controlled action by monkey medial frontal cortex. *Nat. Neurosci.* 10, 240–248. doi: 10.1038/nn1830
- Isoda, M., and Hikosaka, O. (2008). Role for subthalamic nucleus neurons in switching from automatic to controlled eye movement. *J. Neurosci.* 28, 7209–7218. doi: 10.1523/JNEUROSCI.0487-08.2008
- Jenkinson, N., and Brown, P. (2011). New insights into the relationship between dopamine, beta oscillations and motor function. *Trends Neurosci.* 34, 611–618. doi: 10.1016/j.tins.2011.09.003
- Johansen-Berg, H., Gutman, D. A., Behrens, T. E. J., Matthews, P. M., Rushworth, M. F. S., Katz, E., et al. (2008). Anatomical connectivity of the subgenual cingulate region targeted with deep brain stimulation for treatment-resistant depression. *Cereb. Cortex* 18, 1374–1383. doi: 10.1093/cercor/bhm167
- Kajikawa, Y., and Schroeder, C. E. (2011). How local is the local field potential? *Neuron* 72, 847–858. doi: 10.1016/j.neuron.2011.09.029
- Lardeux, S., Pernaud, R., Paleressompoulle, D., and Baunez, C. (2009). Beyond the reward pathway: coding reward magnitude and error in the rat subthalamic nucleus. *J. Neurophysiol.* 102, 2526–2537. doi: 10.1152/jn.91009.2008
- Lauriola, M., Panno, A., Levin, I. P., and Lejuez, C. W. (2013). Individual differences in risky decision making: a meta-analysis of sensation seeking and impulsivity with the balloon analogue risk task. *J. Behav. Decis. Mak.* 27, 20–36. doi: 10.1002/bdm.1784
- Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., et al. (2002). Evaluation of a behavioral measure of risk taking: the Balloon Analogue Risk Task (BART). *J. Exp. Psychol. Appl.* 8, 75–84. doi: 10.1037/1076-898X.8.2.75
- Levy, R., Ashby, P., Hutchison, W. D., Lang, A. E., Lozano, A. M., and Dostrovsky, J. O. (2002). Dependence of subthalamic nucleus oscillations on movement and dopamine in Parkinson's disease. *Brain* 125, 1196–1209. doi: 10.1093/brain/awf128
- Malone, D. A. Jr., Dougherty, D. D., Rezai, A. R., Carpenter, L. L., Friehs, G. M., Eskandar, E. N., et al. (2009). Deep brain stimulation of the ventral capsule/ventral striatum for treatment-resistant depression. *Biol. Psychiatry* 65, 267–275. doi: 10.1016/j.biopsych.2008.08.029
- Nambu, A., Tokuno, H., and Takada, M. (2002). Functional significance of the cortico-subthalamic-pallidal "hyperdirect" pathway. *Neurosci. Res.* 43, 111–117. doi: 10.1016/S0168-0102(02)00027-5
- Nichols, T. E., and Holmes, A. P. (2001). Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum. Brain Mapp.* 15, 1–25. doi: 10.1002/hbm.1058
- Patil, P. G., Carmenta, J. M., Nicolelis, M. A. L., and Turner, D. A. (2004). Ensemble recordings of human subcortical neurons as a source of motor control signals for a brain-machine interface. *Neurosurgery* 55, 27–35; discussion 35–38. doi: 10.1227/01.NEU.0000126872.23715.E5
- Quiroga, R. Q., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput.* 16, 1661–1687. doi: 10.1162/089976604774201631

- Rosin, B., Slovik, M., Mitelman, R., Rivlin-Etzion, M., Haber, S. N., Israel, Z., et al. (2011). Closed-loop deep brain stimulation is superior in ameliorating parkinsonism. *Neuron* 72, 370–384. doi: 10.1016/j.neuron.2011.08.023
- Santaniello, S., Fiengo, G., Glielmo, L., and Grill, W. M. (2011). Closed-loop control of deep brain stimulation: a simulation study. *IEEE Trans. Neural Syst. Rehabil. Eng.* 19, 15–24. doi: 10.1109/TNSRE.2010.2081377
- Tan, H., Zavala, B., Pogosyan, A., Ashkan, K., Zrinzo, L., Foltynie, T., et al. (2014). Human subthalamic nucleus in movement error detection and its evaluation during visuomotor adaptation. *J. Neurosci.* 34, 16744–16754. doi: 10.1523/JNEUROSCI.3414-14.2014
- Tankus, A., Yeshurun, Y., and Fried, I. (2009). An automatic measure for classifying clusters of suspected spikes into single cells versus multiunits. *J. Neural Eng.* 6:056001. doi: 10.1088/1741-2560/6/5/056001
- Whitmer, D., de Solages, C., Hill, B., Yu, H., Henderson, J. M., and Bronte-Stewart, H. (2012). High frequency deep brain stimulation attenuates subthalamic and cortical rhythms in Parkinson's disease. *Front. Hum. Neurosci.* 6:155. doi: 10.3389/fnhum.2012.00155
- Zavala, B. A., Tan, H., Little, S., Ashkan, K., Hariz, M., Foltynie, T., et al. (2014). Midline frontal cortex low-frequency activity drives subthalamic nucleus oscillations during conflict. *J. Neurosci.* 34, 7322–7333. doi: 10.1523/JNEUROSCI.1169-14.2014
- Zavala, B., Brittain, J.-S., Jenkinson, N., Ashkan, K., Foltynie, T., Limousin, P., et al. (2013). Subthalamic nucleus local field potential activity during the Eriksen flanker task reveals a novel role for theta phase during conflict monitoring. *J. Neurosci.* 33, 14758–14766. doi: 10.1523/JNEUROSCI.1036-13.2013
- Zavala, B., Zaghoul, K., and Brown, P. (2015). The subthalamic nucleus, oscillations, and conflict. *Mov. Disord.* 30, 328–338. doi: 10.1002/mds.26072

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Pearson, Hickey, Lad, Platt and Turner. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.