Splitting the Difference

How Does the Brain Code Reward Episodes?

BRIAN KNUTSON AND G. ELLIOTT WIMMER

Department of Psychology, Stanford University, Stanford, California, USA

So, nat'ralists observe, a flea
Hath smaller fleas that on him prey;
And these have smaller still to bite 'em;
And so proceed *ad infinitum.*

—Jonathan Swift, *On Poetry: A Rhapsody*, 1733

ABSTRACT: Animal research and human brain imaging findings suggest that reward processing involves distinct anticipation and outcome phases. Error terms in popular models of reward learning (such as the temporal difference [TD] model) do not distinguish between the updating of expectations in response to reward cues and outcomes. Thus, correlating a single error term with neural activation assumes recruitment of similar neural substrates at each update. Here, we split the error term to separately model reward prediction and prediction errors, and compare the fit of single versus split error terms to functional magnetic resonance imaging (FMRI) data acquired during a monetary incentive delay task. We speculate and find that while the nucleus accumbens computes gain prediction in response to cues, the mesial prefrontal cortex (MPFC) computes gain prediction errors in response to outcomes. In addition to offering a more comprehensive and anatomically situated view of reward processing, split error terms generate novel predictions about psychiatric symptoms and lesion-induced deficits.

KEYWORDS: reward; anticipation; FMRI; human; computation; accumbens; prefrontal

MOTIVATION

Reminiscent of the infinite tower of fleas in Jonathan Swift's sardonic ode to poetry, the image of brain as an ascending hierarchy stands as a lasting

Address for correspondence: Brian Knutson, Department of Psychology, Bldg. 420, Jordan Hall, Stanford, CA 94305. Voice: +650-724-2965; fax: +650-725-5699.
knutson@psych.stanford.edu

Ann. N.Y. Acad. Sci. 1104: 54–69 (2007).
doi: 10.1196/annals.1390.020

contribution of neurologist John Hughlings Jackson's prolific but scattered writings.[1] Based on this hierarchical organization, Jackson predicted that lesions of outermost brain regions should produce not only "negative symptoms" or attenuation of critical faculties, but also "positive symptoms" or accentuation of previously inhibited faculties localized further down. Neurophysiologist Paul MacLean extended the notion of neural hierarchy in his sketch of the "triune brain," by stacking socioemotional concerns atop survival programs, which were in turn crowned by a higher level of symbolic representation.[2] While the triune brain concept has fallen out of favor due to ambiguous specification of subcortical circuitry and the challenge of distinguishing emotional from cognitive function,[3] recent innovations in brain imaging offer new hope for testing hierarchical localization schemes.

Functional magnetic resonance imaging (FMRI) methods have advanced rapidly since the technique's inception in the early 1990s. Conceptual advances in design and analysis soon followed after physical improvements in image acquisition. While initial FMRI experiments were modeled after positron emission tomography (PET) studies, with relatively reduced spatiotemporal resolution (i.e., $\sim$8 mm^3/120 sec), event-related designs have dramatically enhanced FMRI's spatiotemporal resolution (i.e., $\sim$4 mm^3/2 sec), potentially allowing investigators to acquire images at the speed of phenomenology. These new methods call for new models and modes of analysis.

Reward processing represents an evolutionarily conserved, yet environmentally flexible phenomenon that could benefit from temporally precise analysis. To promote survival and reproduction, subjective evaluation often must supersede and direct the processing of other types of information.[4] Thus, flexible evaluators must predict as well as respond to incentive outcomes. If rewarding stimuli are defined as those that an organism will work to obtain, reward processing minimally refers to the unfolding of reward anticipation and outcome phases over time,[5] consistent with a historic ethological distinction between appetitive and consummatory motivation.[6] In prior FMRI research, analyses have suggested that while reward anticipation primarily activates the subcortical nucleus accumbens (NAcc), reward outcomes primarily activate a region of the mesial prefrontal cortex (MPFC).[7] These analyses relied upon simple statistical contrasts that highlighted regions in which one experimental condition elicited greater local oxygen utilization (or "activation") than another. But beyond the critical first step of localization, how can investigators best model the dynamic flow of activity coursing through different brain regions? For instance, when does a signal deviate from baseline, in which direction (up or down), and to what degree? Models that can address these specific questions promise not only the practical benefit of increasing sensitivity to detect activation, but also the theoretical benefit of improving functional understanding of what a given region computes.

Computational models that generate temporally specific predictions have already yielded elegant and profound insights about brain regions implicated

in reward processing. One of the most popular of these models is the temporal difference (TD) model.[8] A key term in TD models is the reward prediction error (i.e., "delta"), or the difference between expected and actual reward. The reward prediction error term can be used to update expectations. For instance, reward outcomes that are initially unexpected create a positive reward prediction error, since such an event exceeds expectations. However, once a cue begins to predict an uncertain reward, the reward prediction error shifts to the unexpected cue onset rather than reward delivery. Thus, like a dog salivating at the sound of a dinner bell, the model learns that specific cues predict eventual rewards.

TD models have recently been used not only to describe behavior, but also the firing rates of midbrain dopamine neurons.[9,10] Specifically, while midbrain dopamine neurons initially increase firing in response to an unexpected juice squirt, after monkeys learn that a cue predicts juice delivery, dopamine firing shifts to cue onset, implying a positive reward prediction error. Importantly, when juice does not follow the cue, dopamine neurons briefly stop firing at the time of expected juice delivery, implying a negative reward prediction error, which rules out surprise or novelty as alternative explanations for changes in firing rate.[10,11] Although the TD error term mimics key features of dopamine neuron firing, it does not indicate whether other brain regions modulate the activity of dopamine neurons, or how they do so.

Based on the temporal distinction between reward anticipation and outcome, as well as a spatial distinction between subcortical and cortical brain regions, we propose an "ascending differences" (AD) split of the TD error term (see TABLE 1). This modification assumes that the brain distinguishes between uncertain and certain events.[12] While uncertain events imply that something may occur in the future (i.e., probability falls between 1 or 0), certain events imply that something has occurred (i.e., probability collapses to either 1 or 0). Thus, uncertainty implies anticipation of an outcome, while certainty implies the outcome itself (which might include either the occurrence or nonoccurrence of an

TABLE 1. AD and TD regressor computation

	Anticipation	Outcome
Temporal difference (TD)		
Gain prediction error (GPE)	CV–EV*	OV–CV
Ascending differences (AD)		
Gain prediction	CV–EV	0
Gain prediction error	0	OV–CV

EV = average(CV) [or average V(t)].
CV = cue magnitude X cue probability [delta(t) at cue or V(t)].
OV = outcome magnitude X outcome probability (i.e., either 0 or 1) [delta(t) at outcome or r(t)].
*Omitting EV from this term does not improve the fit of TD GPE. A future reward prediction term (as found in typical TD error terms) is not included because monetary incentive delay (MID) task trial outcomes are independent (i.e., the current outcome carries no information about future outcomes).

event). Once learning (and thus anticipation) has stabilized, cue presentation elicits a prediction, which can be computed as the cue value minus the baseline expected value (or averaged outcome value up to that point). On the other hand, outcome presentation elicits a prediction error, which can be computed as the outcome value minus the predicted cue value. Thus, the AD modification uses the same error term as TD models, but "splits" reward prediction ("how good might it be?") and reward prediction error ("how good is it?") into two separate terms (see TABLE 1). This split allows for the possibility that different brain regions compute reward prediction and reward prediction error. Here, we name these split terms gain prediction (GP) and gain prediction error (GPE), since the data to be modeled were collected in humans undergoing FMRI as they anticipated and received monetary gains in a monetary incentive delay task.[a] Substitution of the word "gain" for "reward" retains a positive connotation, but allows for gain to be coded relative to some neutral reference point, rather than absolute zero.[13]

Research that initially applied TD models to FMRI data in dynamic tasks involving reward learning revealed that striatal activation correlated with TD reward prediction error,[14,15] a finding borne out by later research.[16,17] After stabilization of learning, however, spatially distinct neural substrates may separately represent GP versus GPEs.[18,19] It is important to extend computational models from dynamic to stable incentive processing tasks, since some psychiatric disorders may involve stable deficits in incentive processing, as suggested by emerging FMRI findings.[20,21] Thus, our goals in this article were to generalize learning models to FMRI data acquired during a stable incentive processing task, and to directly compare the fit of TD versus AD error terms to these data.

MODEL COMPARISON

Based on more than a century of behavioral research,[22] our laboratory has devised a monetary incentive delay (MID) task in which people anticipate and respond to monetary incentives while undergoing FMRI. Monetary incentives provide experimental flexibility because they are nearly universally valued, can be either positive or negative, and can be scaled (features that also facilitate computational modeling). Prior to scanning, subjects are trained on the MID task without pay, and then are shown the cash that they can make while playing in the scanner. In a typical MID task trial, subjects see a cue indicating potential gain or loss of varying magnitudes ($0.00, $0.20, $1.00, $5.00), wait for a brief period (anticipation: 2–3 sec), respond to a rapidly presented target with a

[a]TD models also include a "reward prediction" term (v). However, AD GP as defined here corresponds both with TD (delta(t)) at cue presentation, as well as TD reward prediction (v(t)). These TD terms are collinear in the MID task due to temporal overlap of cues and the short anticipation period that follows. AD GPE corresponds to TD delta(t) at outcome. (see TABLE 1 legend for mappings of AD to TD terms).

button press (~160–260 msec), and receive feedback indicating whether they have either gained or avoided losing money in addition to their cumulative total (outcome: 2 sec), based on the previous cue and whether they pushed the button before the target disappeared.[23] Based on the subject's reaction time, target speed can be adjusted to elicit a desired range of performance (here, an average hit rate of 66%). However, the outcome of each trial is determined independent of the outcomes of previous or subsequent trials.

The analyses presented below focus on changes in brain activation of 26 subjects during gain anticipation (2 sec following cue presentation) and in response to gain outcomes (2 sec during feedback presentation). These data have been previously presented in two reports, but were analyzed with simple unit-weighted contrasts rather than computationally derived regressors.[7,24] To compare TD and AD difference terms, continuous regressors were derived using the formulas described in TABLE 1, with the restriction that each regressor model equals numbers of deviations from baseline, thereby equating the power of each to detect correlated activation. Regressors were convolved with a gamma function to model the lag in hemodynamic response,[25] and entered into otherwise identical multiple regression models (i.e., models also included nuisance regressors that covaried out baseline, linear, and second-order trends for each session, as well as six motion parameter estimates) using Analysis of Functional Neural Images software.[26] Thus, in terms of regressors of interest, the TD model contained gain and loss prediction error regressors, while the AD model contained gain and loss prediction regressors as well as gain and loss prediction error regressors.

Analyses included three stages. First, in localization analyses, group maps were constructed in which model coefficients were tested against the null hypothesis of no activation using *t*-tests ($P < 0.0001$ uncorrected; corrected for the approximate total volume of striatal and mesial frontal gray matter volumes of interest at $P < 0.05$). For the AD model, AD GP and GPE coefficient maps were conjoined to confirm the predicted regional disjunction of activation ($P < 0.0001$ uncorrected). Second, in comparison analyses, coefficients from different models were directly compared using within subjects paired *t*-tests. The key comparisons specifically contrasted AD GP versus TD GPE and AD GPE versus TD GPE, with a focus on NAcc and MPFC volumes of interest ($P < 0.01$ uncorrected). An additional comparison contrasted AD GPE with simple outcome value (i.e., r(t); included in an otherwise identical AD model instead of AD GPE) in the MPFC volume of interest ($P < 0.01$, uncorrected). Third, in verification analyses, activation was averaged and extracted from NAcc and MPFC volumes of interest (based on Knutson *et al.*[7]) and plotted against model predictions for high incentive gain trials (+$5.00 hit and miss, which produce the strongest signal changes).

In localization analyses, statistical maps indicated that TD GPE significantly correlated with NAcc activation and, less robustly, MPFC activation. However, AD regressor maps indicated that while GP maximally correlated with

NAcc activation, GPE instead maximally correlated with MPFC activation (see FIG. 1). For both TD and AD models, loss-related regressors did not positively correlate with activation in these regions (see also TABLE 2). Conjunction of AD GP and AD GPE revealed no conjoint activation

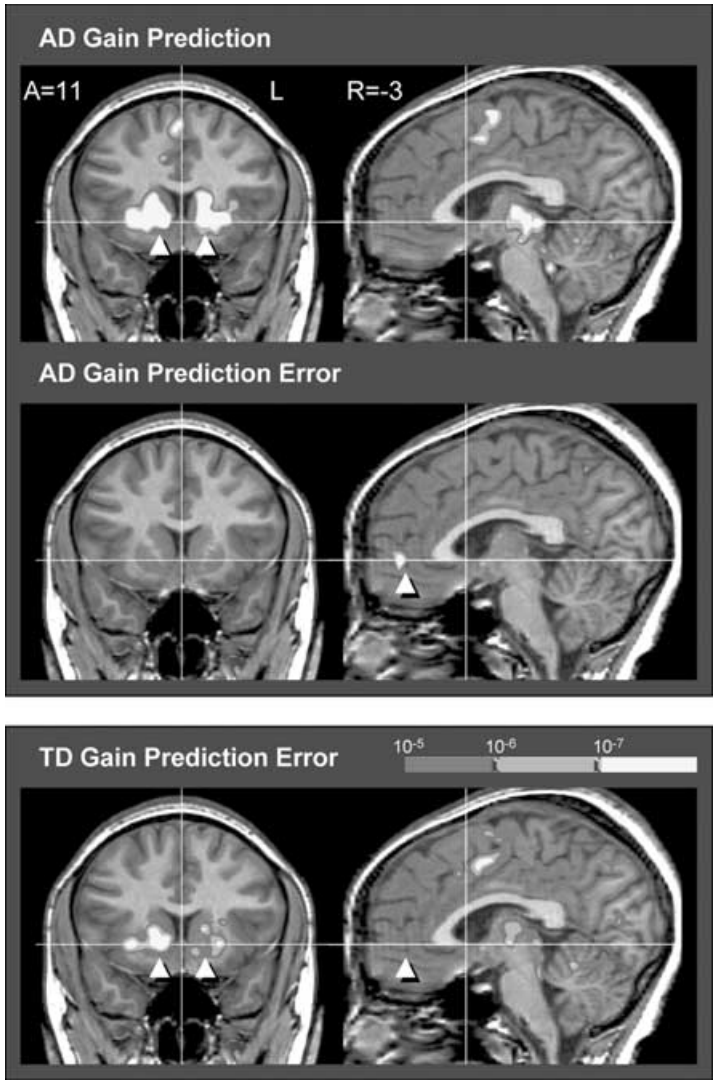

FIGURE 1. Ascending difference (AD) and temporal difference (TD) regressor maps ($n = 26$; $P < 0.0001$, uncorrected).

TABLE 2. AD and TD regressor foci

	Region	Z-score	R	A	S
AD GP (Gain)					
	R anterior cingulate	4.43	9	33	22
		4.83	7	16	36
		5.47	3	7	51
		4.59	4	−13	48
	L middle frontal gyrus	4.27	−31	28	32
	R anterior insula	5.60	27	15	0
	L anterior insula	5.93	−30	19	0
	R nucleus accumbens	6.65	10	7	−1
	L nucleus accumbens	5.85	−11	11	−1
	R caudate	6.08	13	9	10
		5.08	11	1	12
	L caudate	6.50	−12	11	9
		4.61	−7	2	16
	R putamen	6.55	20	7	0
		5.53	27	−5	8
	L putamen	6.23	−18	8	0
		5.98	−25	−2	5
	R thalamus	6.10	6	−20	4
	L thalamus	6.09	−7	−18	4
	R SNc/midbrain	5.36	9	−18	−6
	L SNc/midbrain	4.81	−5	−17	−9
	PAG/midbrain	6.08	0	−24	−7
	R BA 6	4.56	18	2	54
	R SMA/BA 31	4.65	6	−26	49
	R sup. frontal gyrus/BA 6	5.02	25	−10	48
	L sup. frontal gyrus/BA 6	4.63	−17	−7	60
	R precentral gyrus/BA6	4.42	30	−18	48
	L precentral gyrus/BA6	5.18	−33	−15	55
		4.83	−25	−11	63
	R paracentral lobule/BA 6	4.65	6	−25	49
	L postcentral gyrus/BA 3	5.19	−26	−31	58
	R superior parietal lobule	4.93	23	−60	44
	L superior parietal lobule	5.61	−22	−59	56
	R inferior parietal lobule	4.61	27	−48	39
	L inferior parietal lobule	4.66	−29	−48	37
	L precuneus	5.00	−18	−71	19
	R cuneus	4.39	7	−75	7
AD LP (Loss)					
	R superior frontal gyrus	−4.70	2	34	42
		−4.43	4	4	56
	R anterior cingulate	−4.73	6	17	26
	R caudate	−5.34	11	13	9
	L caudate	−4.28	−11	9	11
	R putamen	−5.16	19	10	0
		−5.36	24	1	−2
	L putamen	−4.59	−18	10	0
		−4.50	−23	−4	5

Continued.

TABLE 2. (Continued)

	Region	Z-score	R	A	S
	R thalamus	−5.29	8	−19	12
	L thalamus	−5.28	−5	−17	12
	L precentral gyrus/BA 6	−5.15	−27	−27	50
	R SNc/midbrain	−4.59	6	−18	−5
	PAG/midbrain	−4.89	2	−26	−16
AD GPE (Gain)					
	R MPFC	4.49	4	50	0
	L MPFC	5.50	−4	50	0
	L middle frontal gyrus	4.34	−27	29	37
	R NAcc	4.34	12	12	−7
	Posterior cingulate	4.83	0	−37	30
	R posterior cingulate	4.78	7	−52	15
	L posterior cingulate	4.74	−3	−57	15
	R paracentral lobule	4.31	2	−38	54
	R cuneus	5.10	25	−87	9
AD LPE (Loss)	(No regions survive threshold and cluster criteria*)				
TD GPE (Gain)					
	R anterior cingulate	4.42	4	25	30
		5.64	7	3	46
		4.89	2	−3	31
	L anterior cingulate	5.56	−4	1	44
		4.44	−4	−15	45
	R cingulate	4.59	6	−23	49
	L cingulate	4.30	−8	−13	38
	R superior frontal gyrus	4.12	7	44	21
	L superior frontal gyrus	4.26	−10	34	20
	L genual cingulate	5.29	−7	41	−3
	R middle frontal gyrus	4.56	30	−3	56
		4.63	22	0	47
	R anterior insula	5.52	31	16	1
	L anterior insula	4.34	−33	18	2
	R nucleus accumbens	5.17	11	10	−4
	L nucleus accumbens	4.68	−10	12	−2
	R caudate	5.51	13	8	10
		5.03	11	−1	−15
	L caudate	4.50	−14	10	10
		4.52	−15	1	19
	R putamen	5.50	17	10	0
		5.87	23	−1	4
	L putamen	5.04	−18	7	0
	R thalamus	5.88	8	−18	4
	L thalamus	5.01	−7	−12	6
	R SNc/midbrain	5.15	3	−15	−7
	PAG/midbrain	4.97	−3	−30	−13
	R precentral gyrus/BA6	4.10	30	−22	56

Continued

TABLE 2. (Continued)

	Region	Z-score	R	A	S
	L precentral gyrus/BA6	4.26	−30	−9	57
	R precuneus	4.27	12	−42	50
		4.55	8	−68	23
	R inferior parietal lobule	4.60	22	−59	41
	L inferior parietal lobule	4.74	−30	−54	42
		5.34	−25	−62	37
		4.74	−12	−68	42
	R cuneus	4.88	11	−60	9
TD LPE (loss)	No regions survive threshold and cluster criteria*				

*$P < 0.0001$, uncorrected; cluster = 3 voxels.

in the NAcc or MPFC (or any other region). Relaxing the cluster criterion revealed a single conjointly activated voxel in the NAcc, but this region was more than two orders of magnitude smaller (50 mm^3) than the region activated by AD GP alone (>15,000 mm^3). Comparison of AD GPE with simple outcome value revealed that AD GPE more robustly correlated with MPFC activation than outcome value (left TC: –4,50,–3; Z = 3.02; right TC: 4,50,3; Z = 3.51; ps < 0.01). This result is consistent with MPFC activation time course plots (see FIG. 2), which indicate not only that gain outcomes increase MPFC activation (predicted by both AD GPE and outcome), but also that nongain outcomes decrease MPFC activation (predicted only by AD GPE).

In comparison analyses, AD GP correlated with NAcc activation more robustly than TD GPE, while AD GPE correlated with MPFC activation marginally more than TD GPE, as predicted ($P < 0.01$, warm colors, see FIG. 3). Conversely, and consistent with less spatial specificity for the TD error term, TD GPE correlated with MPFC activation more robustly than AD GP, while TD GPE correlated with striatal activation more robustly than AD GPE (particularly in the lateral putamen and dorsal caudate; cool colors; see FIG. 3 and TABLE 3). Intriguingly, TD GPE correlated more robustly with anterior cingulate activation than both AD GP and AD GPE, implying that the TD GPE term may more closely model activation in regions other than the NAcc or MPFC.

In verification analyses, visual inspection of activation time course plots for high gain (+$5.00) conditions (hits and misses) confirmed findings from direct comparisons. Specifically, AD GP fit NAcc time course data more closely than did TD GPE, while AD GPE fit MPFC time course data more closely than did TD GPE (see FIG. 2). The AD model fits appeared closest for hits, with miss data falling somewhere between the predictions of AD and TD error terms in the NAcc (see TABLE 3 for direct statistical contrasts of models—plots provide visual comparison only).

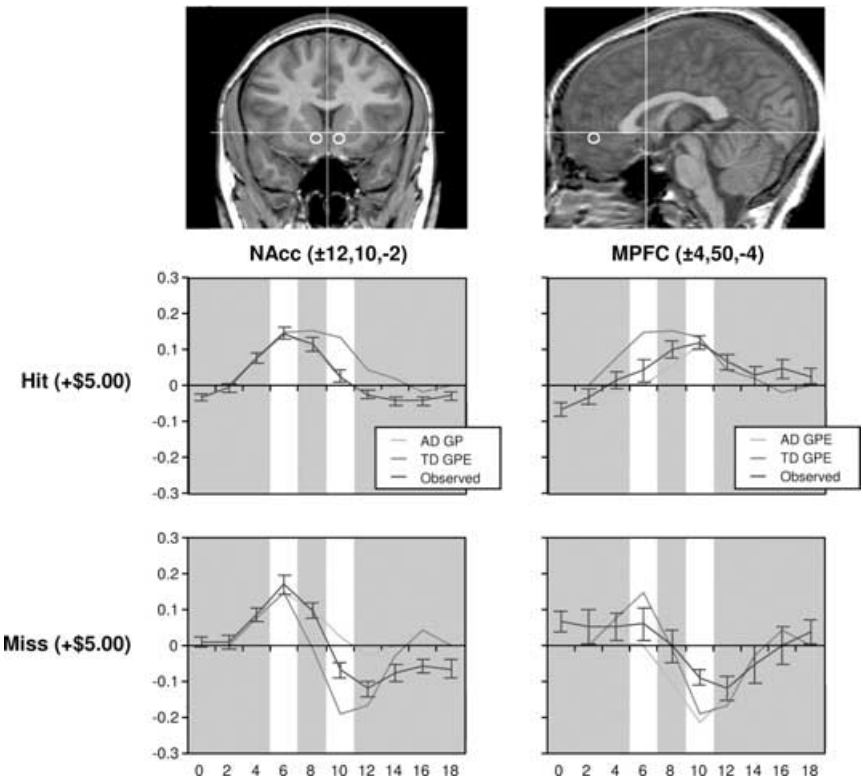

FIGURE 2. Relation of AD and TD regressor predictions to NAcc and MPFC activation time course data for +$5.00 hit and miss trials (mean ± SEM; $n = 26$; white bars indicate predicted peak activation for anticipation and outcome phases of trials lagged by 6 sec).

As suspected, AD error terms highlighted a dissociation undetected by the TD error term: while NAcc activation correlated with AD GP after cue presentation, MPFC activation correlated instead with AD GPE in response to outcomes. Both direct within-subject comparisons and visualization of activation time courses confirmed this dissociation. While the present AD split of the TD error term is conceptually simple and easy to implement, it represents more of a beginning than an end. Since the AD error terms are optimized for the MID task, which is stable and involves minimal learning, the present AD modification might require further elaboration to generalize to more dynamic learning scenarios. For instance, a parameter representing memory decay over time could be added to the GPE term. Also, rather than computing baseline expected value as an average, baseline expected value could be biased toward recent experiences. Finally, outcomes might also elicit uncertainty (i.e., which might serve as both outcomes and cues for future rewards) when they

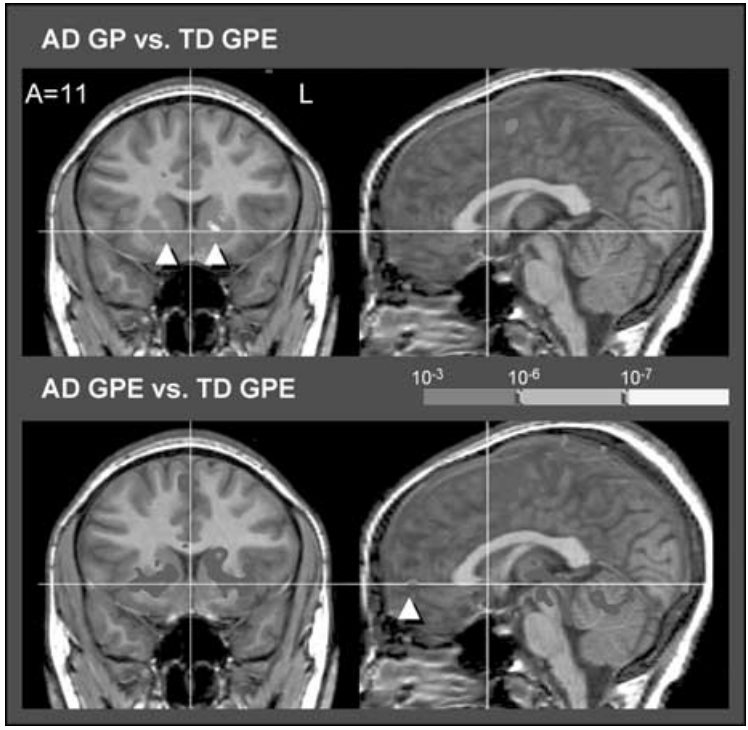

FIGURE 3. AD versus TD regressor comparison ($n = 26$, within-subjects; warm colors indicate better fit for AD terms while cool colors indicate better fit for TD terms; $P < 0.01$, uncorrected).

substantially deviate from the range of predictions suggested by a prior cue, or when they are correlated with outcomes of neighboring trials. Nonetheless, even the present simple AD split yields novel and testable predictions that extend beyond those generated by a single TD error term.

IMPLICATIONS

In the context of a stable incentive processing task, a temporal differences reward prediction error term correlated with activation in both NAcc and, to a lesser extent, MPFC regions, replicating a pattern of findings reported in a growing number of reports.[14,19,27,28] However, the currently proposed split of the TD error term revealed a spatiotemporal dissociation in which GP correlated most closely with NAcc activation, while GPE correlated more closely with MPFC activation.

In each region, AD GP and GPE terms predicted both increases and decreases in activation driven by local blood oxygenation. While changes in oxygenation correlate with changes in postsynaptic neural activity, researchers still have not determined how the two are physiologically linked.[29] Dopamine neurons have been consistently implicated in reward processing, and project from the ventral tegmental area (VTA) of the midbrain to both the NAcc and MPFC. Electrophysiological studies show that these dopamine neurons fire at an average rate of approximately 5 impulses per second,[30] and this firing temporarily increases when animals anticipate rewards or receive unexpected rewards, but temporarily decreases when animals fail to receive expected rewards.[10] Further, monkey research suggests that injection of dopamine-releasing agents, such as amphetamine, can increase FMRI activation in the NAcc, while concurrent dopamine depletion abolishes these amphetamine-induced increases.[31] Thus, dopaminergic modulation of postsynaptic targets may contribute to changes in activation visualized with FMRI.[32] However, the independence of signal changes in the NAcc and MPFC suggest that activation in these regions is not modulated solely by dopamine release, raising the possibility that NAcc and MPFC may exert differential control on VTA dopamine neurons via descending projections.[33] While not discussed here in detail (due to potential artifactual warping of midbrain regions), VTA activation more closely correlated with AD GP than GPE (see TABLE 1), a pattern that has been replicated in other FMRI studies.[34,35] Future studies with enhanced temporal resolution may better elucidate the influence of NAcc and MPFC activation on VTA activity.

The AD split of the TD error term offers both practical and theoretical benefits as a tool for predicting changes in FMRI activation. From a practical standpoint, models that enhance sensitivity for detecting activation promise to save investigators both time and money. These models may also help investigators to decompose and better understand symptoms related to psychiatric disorders. To have clinical relevance, a computational model must not only predict brain activity, but the predicted brain activity must then correlate with behavioral phenomena (e.g., affect, behavior, cognition) of psychiatric importance. In stable incentive processing tasks, NAcc activation has been correlated with the experience of positive but not negative aroused affect in healthy individuals undergoing both FMRI[24] and PET studies.[36] Deficits in positive arousal have been documented in psychiatric disorders ranging from affective disorders to addiction to schizophrenia. In the context of the MID task, we have recently observed that schizophrenics (both never-medicated and treated with traditional neuroleptics) show specific deficits in NAcc activation during GP. Further, the extent of this blunting correlates with a chronic absence of positive arousal (called "negative symptoms" in the psychiatric literature).[20,37] Thus, some schizophrenics may suffer from a deficit in GP, which in turn manifests as negative symptoms. A model that conflates GP and GPE would not have elucidated this specific deficit. Thus, this set of findings exemplifies

TABLE 3. AD versus TD regressor direct comparison foci ($P<.01$, uncorrected; cluster = 3 voxels; gain regressors only)

	Region	Z-score	R	A	S
AD GP vs. TD GPE**					
AD > TD	R cingulate	3.07	7	21	42
		2.86	18	3	53
	L medial frontal gyrus	3.74	−3	2	54
	R anterior insula	3.52	32	17	7
	L anterior insula	2.83	−31	15	6
	R NAcc/putamen	4.80	12	12	0
	L putamen	4.28	−18	14	−5
		3.35	−25	5	5
	R thalamus	3.29	15	−22	10
	L thalamus	3.35	−18	−15	2
	R SNc/midbrain	2.95	4	−17	−3
	L SNc/midbrain	3.26	−2	−15	−2
	R precentral gyrus/BA 6	3.95	27	−11	50
	L precentral gyrus/BA 6	3.05	−29	−16	60
		3.51	−33	−23	46
TD > AD	R MPFC	−3.88	3	56	8
		−3.44	3	49	0
		−3.84	6	48	25
	L superior frontal gyrus	−4.30	−15	36	39
	R subgenual cingulate	−2.85	2	18	−10
	R posterior cingulate	−4.15	10	−53	15
	L posterior cingulate	−3.86	−3	−50	33
	R parahippocampal gyrus	−3.23	12	−7	−15
	L parahippocampal gyrus	−3.89	−18	−15	−15
AD GPE vs. TD GPE					
AD > TD	MPFC	3.00	0	53	2
	R posterior cingulate	3.36	7	−53	30
		2.89	7	−53	16
	R insula	3.53	41	−12	15
TD > AD	R cingulate	−3.14	7	15	37
		−3.98	5	10	52
	L cingulate	−3.28	−4	4	45
		−4.22	−4	−16	49
	R anterior insula	−4.37	24	19	−1
	L anterior insula	−4.64	−30	19	−3
	L caudate/ putamen	−4.76	−16	9	4
	R caudate	−4.09	11	17	7
		−3.14	14	5	15
	L caudate	−4.60	−16	11	15
		−3.44	−12	3	15
	R putamen	−4.14	16	12	−1
		−3.99	22	1	3
	L putamen	−3.79	−20	11	−1
		−3.36	−22	−7	8
	R thalamus	−4.17	5	−16	6
	L thalamus	−4.35	−6	−16	5

Continued.

TABLE 3. (Continued)

	Region	Z-score	R	A	S
AD GP vs. TD GPE**					
	PAG/midbrain	−3.30	−4	−31	−11
	R SNc/midbrain	−3.19	3	−18	−7
	L SNc/midbrain	−3.80	−7	−16	−4
	R precentral gyrus/BA 6	−3.34	15	−18	64
	L precentral gyrus/BA 6	−3.59	−33	−19	58
	R superior parietal lobule	−3.42	25	−49	54

**3 neighboring voxels at $P < 0.01$.

how properly specified computational models may help investigators to better probe covert phenomena (e.g., affective experience) that nonetheless can have an overt impact (e.g., on feelings, thoughts, or behavior) on psychiatric health.

From a theoretical standpoint, computational models must not only correlate with brain activation, but should also approximate the operation of neural mechanisms. The selection of models that fit neural constraints can thus be framed as a continuing journey of closer approximations. The AD split augments the TD error term by suggesting that the brain distinguishes between uncertain anticipation and certain outcomes, as well as between gain and loss.[38] Of course, such distinctions raise new questions about which neural circuits support the computation of different error terms, and how information from these circuits then combines to coherently inform learning and behavior. Eventual answers to such questions may help investigators to better isolate distinct neuropsychological components that contribute to complex disorders.

Distinct spatial correlates of AD GP and GPE imply that output from lower and higher centers may combine, either cooperatively or competitively, to produce a given behavior. For instance, if the NAcc computes GP while the MPFC computes GPE, MPFC lesions might severely impair relearning of reward associations, but paradoxically spare or even accentuate preexisting reward associations.[5] How and where these terms combine to channel behavior remains to be discovered. But as Jackson prophesied long ago, optimal function may require dynamic coordination of brain circuits both high and low.

ACKNOWLEDGMENTS

We thank Jamil Bhanji for technical assistance, as well as Jeffrey C. Cooper, Matthew T. Kaufman, Paul King, and an anonymous reviewer for comments on previous drafts of this manuscript. During manuscript preparation, BK was supported by a National Alliance for Research on Schizophrenia and Depression Young Investigator Award and NIDA Grant DA020615-01.

REFERENCES

1. TAYLOR, J. 1958. Selected Writings of John Hughlings Jackson. Basic Books. New York.
2. MACLEAN, P.D. 1990. The Triune Brain in Evolution. Plenum Press. New York.
3. LEDOUX, J.E. 1993. Emotional memory: in search of systems and synapses. *In* Brain Mechanisms: Papers in Memory of Robert Thompson, Vol. 702. F.M. Crinella & J. Yu, Eds.: 149–157. Ann. N.Y. Acad. Sci. New York.
4. PANKSEPP, J. 1998. Affective Neuroscience: The Foundations of Human and Animal Emotions. Oxford University Press. New York.
5. KNUTSON, B. & J.C. COOPER. 2005. Functional magnetic resonance imaging of reward prediction. Curr. Opin. Neurol. **18:** 411–417.
6. CRAIG, W. 1918. Appetites and aversions as constituents of instincts. Biol. Bull. **34:** 91–107.
7. KNUTSON, B. *et al*. 2003. A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: characterization with rapid event-related FMRI. Neuroimage **18:** 263–272.
8. SUTTON, R.S. & A.G. BARTO. 1990. Time-derivative models of Pavlovian reinforcement. *In* Learning and Computational Neuroscience: Foundations of Adaptive Networks. M. Gabriel & J. Moore, Eds.: 497–537. MIT Press. Boston.
9. MONTAGUE, P.R., P. DAYAN & T.J. SEJNOWSKI. 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. J. Neurosci. **16:** 1936–1947.
10. SCHULTZ, W., P. DAYAN & P.R. MONTAGUE. 1997. A neural substrate of prediction and reward. Science **275:** 1593–1599.
11. MIRENOWICZ, J. & W. SCHULTZ. 1996. Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. Nature **379:** 449–451.
12. FIORILLO, C.D., P.N. TOBLER & W. SCHULTZ. 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. Science **299:** 1898–1902.
13. KAHNEMAN, D. & A. TVERSKY. 1979. Prospect theory: an analysis of decision under risk. Econometrica **47:** 263–291.
14. MCCLURE, S.M., G.S. BERNS & P.R. MONTAGUE. 2003. Temporal prediction errors in a passive learning task activate human striatum. Neuron **38:** 338–346.
15. O'DOHERTY, J.P. *et al*. 2003. Temporal difference models and reward-related learning in the human brain. Neuron **38:** 329–337.
16. KIM, H., S. SHIMOJO & J.P. O'DOHERTY. 2006. Is avoiding an aversive outcome rewarding? Neural substrates of avoidance learning in the human brain. Public Library of Science Biology **4:** e233.
17. DAW, N.D. & K. DOYA. 2006. The computational neurobiology of learning and reward. Curr. Opin. Neurobiol. **16:** 199–204.
18. KNUTSON, B. *et al*. 2001. Dissociation of reward anticipation and outcome with event-related FMRI. Neuroreport **12:** 3683–3687.
19. HARUNO, M. *et al*. 2004. A neural correlate of reward-based behavioral learning in caudate nucleus: a functional magnetic resonance imaging study of a stochastic decision task. J. Neurosci. **24:** 1660–1665.
20. JUCKEL, G. *et al*. 2006. Dysfunction of ventral striatal reward prediction in schizophrenia. Neuroimage **29:** 409–416.
21. SCHERES, A. *et al*. 2007. Ventral striatal hyporesponsiveness during reward anticipation in Attention Deficit/Hyperactivity Disorder. Biol. Psychiatry **61:** 720–724.

22. PAVLOV, I.P. 1927. Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex. Oxford University Press. London.
23. KNUTSON, B. *et al*. 2000. FMRI visualization of brain activity during a monetary incentive delay task. Neuroimage **12:** 20–27.
24. KNUTSON, B. *et al*. 2001. Anticipation of increasing monetary reward selectively recruits nucleus accumbens. J. Neurosci. **21:** RC159.
25. COHEN, M.S. 1997. Parametric analysis of fMRI data using linear systems methods. Neuroimage **6:** 93–103.
26. COX, R.W. 1996. AFNI: software for analysis and visualization of functional magnetic resonance images. Comp. Biomed. Res. **29:** 162–173.
27. GOTTFRIED, J.A., J. O'DOHERTY & R.J. DOLAN. 2003. Appetitive and aversive olfactory learning in humans studied using event-related functional magnetic resonance imaging. J. Neurosci. **22:** 10829–10837.
28. O'DOHERTY, J. *et al*. 2004. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science **304:** 452–454.
29. LOGOTHETIS, N.K. & B.A. WANDELL. 2004. Interpreting the BOLD signal. Ann. Rev. Physiol. **66:** 735–769.
30. WIGHTMAN, R.M. & D.L. ROBINSON. 2002. Transient changes in mesolimbic dopamine and their association with 'reward.' J. Neurochem. **82:** 721–735.
31. JENKINS, B.G. *et al*. 2004. Mapping dopamine function in primates using pharmacological magnetic resonance imaging. J. Neurosci. **24:** 9553–9560.
32. KNUTSON, B. & S.E.B. GIBBS. 2007. Linking nucleus accumbens dopamine and blood oxygenation. Psychopharmacology **191:** 813–822.
33. SHIZGAL, P. 1997. Neural basis of utility estimation. Curr. Opin. Neurobiol. **7:** 198–208.
34. KNUTSON, B. *et al*. 2005. Distributed neural representation of expected value. J. Neurosci **25:** 4806–4812.
35. ADCOCK, R.A. *et al*. 2006. Reward-motivated learning: mesolimbic activation precedes memory formation. Neuron **50:** 507–517.
36. DREVETS, W.C. *et al*. 2001. Amphetamine-induced dopamine release in human ventral striatum correlates with euphoria. Biol. Psychiatry **49:** 81–96.
37. JUCKEL, G. *et al*. 2006. Dysfunction of ventral striatal reward prediction in schizophrenic patients treated with typical but not atypical neuroleptics. Psychopharmacology **187:** 222–228.
38. YACUBIAN, J. *et al*. 2006. Dissociable systems for gain-and loss-related value predictions and errors of prediction in the human brain. J. Neurosci. **26:** 9530–9537.