Predicting risky choices from brain activity patterns

Sarah M. Helfinstein[a,1], Tom Schonberg[a], Eliza Congdon[b,c], Katherine H. Karlsgodt[d,e], Jeanette A. Mumford[f], Fred W. Sabb[b,c,g], Tyrone D. Cannon[h,i], Edythe D. London[b,c,g,j], Robert M. Bilder[b,c,g,k], and Russell A. Poldrack[a,f,l]

[a]Imaging Research Center, and Departments of [f]Psychology and [l]Neuroscience, University of Texas at Austin, Austin, TX 78712; Departments of [b]Psychiatry and Biobehavioral Sciences, [j]Molecular and Medical Pharmacology, and [k]Psychology, [g]Brain Research Institute, and [c]Semel Institute, University of California, Los Angeles, CA 90095; [d]Department of Psychiatry, Zucker Hillside Hospital, Queens, NY 11004; [e]Feinstein Institute for Medical Research, Manhasset, NY 11030; and Departments of [h]Psychology and [i]Psychiatry, Yale University, New Haven, CT 06520

Previous research has implicated a large network of brain regions in the processing of risk during decision making. However, it has not yet been determined if activity in these regions is predictive of choices on future risky decisions. Here, we examined functional MRI data from a large sample of healthy subjects performing a naturalistic risk-taking task and used a classification analysis approach to predict whether individuals would choose risky or safe options on upcoming trials. We were able to predict choice category successfully in 71.8% of cases. Searchlight analysis revealed a network of brain regions where activity patterns were reliably predictive of subsequent risk-taking behavior, including a number of regions known to play a role in control processes. Searchlights with significant predictive accuracy were primarily located in regions more active when preparing to avoid a risk than when preparing to engage in one, suggesting that risk taking may be due, in part, to a failure of the control systems necessary to initiate a safe choice. Additional analyses revealed that subject choice can be successfully predicted with minimal decrements in accuracy using highly condensed data, suggesting that information relevant for risky choice behavior is encoded in coarse global patterns of activation as well as within highly local activation within searchlights.

fMRI | machine learning | decision-making

In daily life, we are regularly confronted with decisions that require choosing between a risky option and a safe one. The safe option typically has a single known outcome, whereas the risky option has multiple possible outcomes, some positive and some negative. Some of these decisions, such as choosing whether to drive while intoxicated or experiment with drugs, can have serious consequences for health and well-being, yet it is poorly understood what leads us to choose these risky courses of action over their safer alternatives.

Although many studies have examined the neural correlates of risk taking (1–4), most, but not all (5, 6), have used experimental paradigms that utilize the economic definition of risk as variance in outcomes. As has been observed elsewhere (7), lay people generally conceive of risk not as variance in outcomes but as exposure to negative outcomes, and real-world risk-taking behavior is only weakly related to performance on variance-based risk-taking tasks (7, 8).

To examine the cognitive processes at work while making risky decisions in a way that relates to real-world risk taking, we used the Balloon Analog Risk Task (BART) (9), in which subjects receive points as they pump up balloons but risk losing those points should the balloon explode before they choose to stop pumping and "cash out." Each pump opportunity is a risky decision, where subjects must choose whether to pump again to gain more points or to cash out to secure those points already accrued. The structure of the task, where subjects make sequential risky choices with feedback, is common to many real-world risk-taking situations (10) and matches both the economic and lay definition of risk, in that each successive pump opportunity for a given balloon has both greater variance in possible outcomes and increased exposure to loss. Performance on this

task has also been shown, in numerous behavioral studies, to relate to self-reported sensation seeking (9, 11) and to naturalistic risk-taking behaviors, such as smoking, drug use, sexual risk-taking, and unsafe driving behaviors (9, 11–14). Because performance on this task consistently correlates with naturalistic risk-taking behaviors, the cognitive processes at work during the task are likely to be comparable to those used during real-world risky decision making.

In addition to the focus on the neural correlates of economic risk, much previous research has emphasized how brain activity changes when levels of risk or related constructs, such as ambiguity, uncertainty, or variance, are manipulated. A large network of regions that are sensitive to risk has been identified, including the insula, frontal cortices, striatum, cingulate cortex, thalamus, and parietal lobes (1–6, 15–18). However, these studies have left a related question largely unaddressed: When confronted with a decision with a given amount of risk, what key differences in neural activity predict whether the subject will go on to make the risky or safe choice? Identifying the regions that differentiate between these two outcomes when the objective amount of risk is the same is critical to understanding risky decision making, because these regions are the ones where the information encoded is most likely to influence the decision-making process. By examining the functionality of these regions, we can then infer the types of information and cognitive processes most likely to shape the outcome of the risky decision. Previous studies (15, 19–23) have examined differences in activation on trials where subjects made either risky or safe choices; however, due to the tight temporal coupling between the making of the decision and the subsequent anticipation of its outcome, it is difficult to determine whether those differences precede and influence the decision or, instead, reflect differences in anticipation of the outcomes of risky or safe choices that occur after the decision has been made.

Significance

Previous studies have examined the neural correlates of risk, but it is unknown if patterns of brain activity can predict choices in risky decision-making. We used functional MRI data to predict choice behavior in subjects while they performed a naturalistic risk-taking task. We found choices on subsequent trials could be predicted with high accuracy when condensing each individual's brain activity to two values, indicating that choice behavior is encoded even in coarse activation patterns. A searchlight analysis demonstrated that choices can also be predicted based on localized activity patterns within neural networks involved in cognitive control. These regions show greater activation prior to safe choices than risky choices, suggesting that control systems play a key role in inhibiting risky choices.

To determine the brain regions where neural activity is directly related to the subsequent choice to take or abstain from risk, we performed a classification analysis on functional MRI data from trials immediately preceding either risky or safe choices to determine whether a machine-learning algorithm could use the data successfully to learn to predict the category to which a given data sample belonged. We used data from trials preceding either risky or safe choices, rather than data from risky or safe choice trials themselves, because trials preceding both risky and safe choices were always successful risky choices due to the structure of the BART task. Because subjects made identical choices and received identical outcomes on trials preceding both risky and safe choices, using data from these trials ensured that any effects of motoric response or anticipation of outcome would be identical for both conditions, allowing us to determine which brain regions encode information that is predictive of the upcoming decision. To ensure that any predictive validity held across individuals, we collected data from a large number of subjects and used a between-subject cross-validation strategy in which safe and risky choice trials were matched on the amount of risk and potential loss. Our goal was first to establish whether neural encoding on trials preceding either a risky or safe choice was sufficiently distinctive that it could be used to predict subsequent subject behavior. After establishing this was the case, our second goal was to determine which regions of the brain encoded information that was most predictive of choice, using a searchlight analysis. Third, we wanted to establish whether prediction relied exclusively on information from local clusters of highly informative voxels that could be identified in searchlight analyses or if a coarse encoding of information across large portions of the brain was also sufficient to predict subject choice. To test this, we extracted mean activation from regions that were more active before a cash out than before a pump and mean activation from regions more active before a pump than before a cash out, and used these two averaged values alone to predict subject choice.

Results

Behavioral Data. One hundred eight participants completed an average of 18.03 experimental balloons (SD = 3.02) and, of those, cashed out on 11.01 (SD = 3.44). They pumped an average of 3.889 times (SD = 0.958) per balloon across all experimental balloons and 4.438 times (SD = 1.375) on balloons that ended in a cash out. Only those pump events preceding a cash out that could be matched with a pump event preceding another pump at the same level of risk were included in the classification analyses, yielding an average of 7.519 (SD = 2.797) useable pre-cash out events per subject. For those events included, reaction times were 689.2 ms (SD = 547.6 ms) for pre-cash out events and 546.3 ms (SD = 252.9 ms) for prepump events (significant difference: $t_{107} = 3.08$, $P = 0.003$).

Functional MRI Data. Our first objective was to identify whether the brain states preceding risky and safe choices were sufficiently unique to be discriminable. To that end, we trained a support vector machine classification algorithm using sixfold cross-validation to classify between mean activation maps for trials that immediately preceded a decision to cash out (i.e., safe choices) vs. maps for trials that immediately preceded a decision to continue inflating the balloon (i.e., risky choices) (Fig. 1). Cross-validation was performed across individuals; the classifier was trained on five-sixths of the subjects (90 of 108) and then applied to the remaining subjects to discriminate between trials that preceded risky vs. safe choices. To ensure against variability due to fold assignment, this was performed 50 times with random fold assignment and the results were averaged. We used trials that immediately preceded the decision rather than the decision-making trial itself to ensure that the classifier would not be able to discriminate between trials based on motor or reward effects,
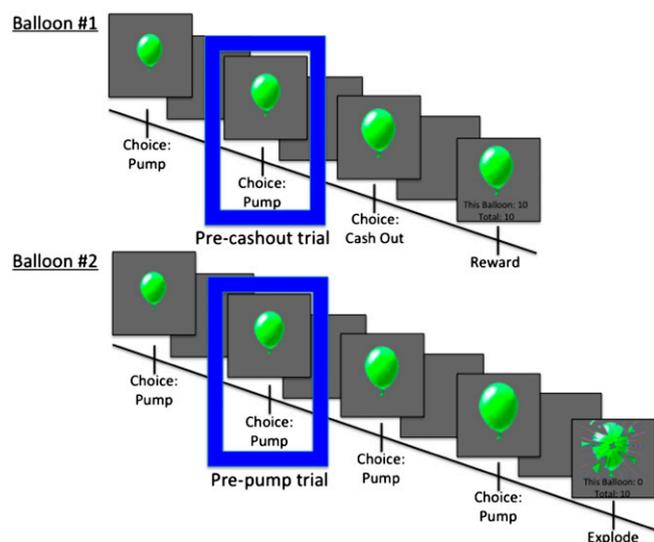

Fig. 1. Sample trials from the BART task. On each balloon, subjects made successive choices about whether to pump or cash out. They earned five points for each pump made before cashing out; however, if the balloon exploded before they cashed out, no points were earned for that balloon. Thus, each pump opportunity represented a risky decision where the subject could choose a certain, safe option (cash out) or an uncertain, risky option (pump). Trials boxed in blue indicate trials suitable for classification analysis; trials are matched for level of risk and subject choice on the trial but differ on subject choice on the subsequent trial. Note that although many "pre-pump" trials are present in the figure, only the boxed prepump trial was selected because it has the same level of risk as the paired "pre-cash out" trial.

which differ at the time of the actual decision. Trials in each category were matched for the number of pumps the subject had made on that balloon before the trial of interest, so that all variables were equated across choice conditions. Prepump and pre-cash out events were modeled with variable epoch regressors to control for effects of reaction time differences between the two conditions (24). Thus, any differences seen between trials preceding risky and safe choices could only be attributed to subjects' cognitive processing on these trials, rather than to differences in the objective features of the trials themselves or the length of time spent engaging in cognitive processing.

Overall, the classifier was effective at discriminating between maps associated with risky (preceding a pump) and safe (preceding a cash out) trials: Using a whole-brain classifier, we were able to classify activation maps correctly as either risky or safe 71.8% of the time. The statistical significance of this accuracy was assessed nonparametrically by randomizing labels within subjects and performing the analysis 500 times to obtain an empirical null distribution. The maximum accuracy across all runs included in the null distribution was 67%; thus, the observed accuracy of the classification was significant at $P < 0.002$.

Next, we identified the regions of the brain where activity was most predictive of risky vs. safe choices. To identify these regions, we conducted a searchlight analysis (25), which revealed several regions where classification ability was significantly better than chance, including bilateral parietal and motor regions, anterior cingulate cortex, bilateral insula, and bilateral lateral orbito-frontal cortex (OFC) (Fig. 2 and Table 1). To establish the cognitive functions in which these regions were most involved, we conducted a formal reverse inference analysis using NeuroSynth (26), correlating our searchlight map with the neural activation maps for each term in the NeuroSynth database. Results (Table 2) revealed that the regions identified in searchlight analyses are involved in control processes, indicating that control networks function in systematically different ways before risky vs. safe choices.
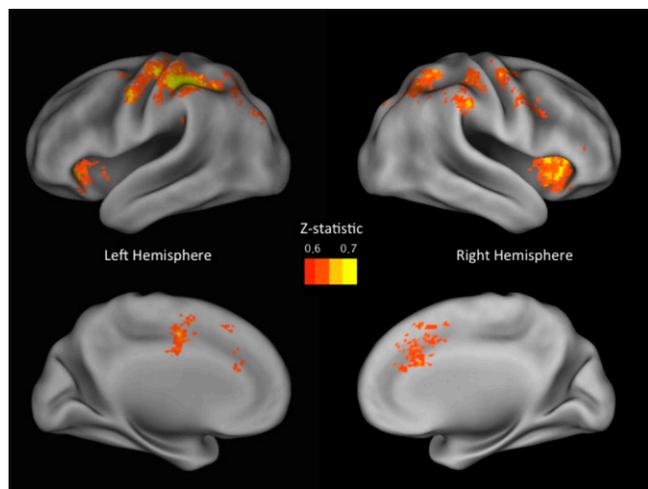
Fig. 2. Classification searchlight analysis of pre-cash out vs. prepump trials. Colored voxels indicate centers of searchlights where the classifier could successfully discriminate between prepump and pre-cash out activation patterns (searchlight classification >60%, whole-brain cluster-corrected $P <$ 0.05 via comparison with 1,000 random permutations). (Scale: 60–70%.) Activation maps were projected onto an inflated average cortical surface of the Population-Average, Landmark-, and Surface-Based (PALS) atlas using the multifiducial mapping technique of Van Essen (29).

In addition to the local sensitivity identified in the searchlight analysis, we examined whether choices could reflect a larger scale balance between systems that show increased activation before the choice of a risky vs. safe option. To assess this, we used a cross-validated feature selection approach to determine whether mean activation across all voxels that were active before either risky or safe choices was sufficient to classify trials successfully as either risky or safe. To do this, using the same design matrix that served as the basis for the classification analysis, we calculated a univariate contrast map for activation preceding a risky choice vs. activation preceding a safe choice, using a per-voxel threshold of $t > 2$ (risky > safe) or $t < -2$ (safe > risky). Only training data were used to generate these contrast maps to avoid any "peeking" bias. We then calculated the mean activation across the "subsequent risky choice" and "subsequent safe choice" voxels for each subject in the training set. This yielded two mean values for each participant, which were used to train a classifier. Using just these two derived means from the same two sets of voxels, it was possible to classify subsequent choices in the test set with above-chance accuracy (67.0%; $P < 0.002$ via comparison with 500 random permutations); these data are plotted in Fig. 3. The logistic regression-based classification boundary indicates that mean activation in both "risky choice"

and "safe choice" regions provides unique information that aids in classification. We confirmed this by performing classifications using only one mean at a time. We found that classification using only the mean of risky > safe regions (56.7%) or the mean of safe > risky regions (59.0%) was significantly worse than the joint classification (risky > safe vs. both: $t_{98} = 58.2$, $P < 0.001$; safe > risky vs. both: $t_{98} = 45.0$, $P < 0.001$). These results indicate that even coarse information about levels of activation in key networks can be used successfully to predict subsequent choice behavior in a risky decision-making task, and that overall levels of activity in systems sensitive to both risk seeking and risk avoidance contribute separately to the ability to predict choice. Together with the searchlight results, these findings suggest that the neural computations underlying risk are reflected in regional signals in parietal, insular, and medial frontal cortices, and in the global balance of activity between regions associated with risk seeking and those associated with risk avoidance.

To determine whether the local information identified in the searchlight analysis contributed to mean activation in the regions sensitive to risk seeking, risk avoidance, or both, we examined the overlap between searchlight maps and the contrast maps used in the coarse-coding analysis. We found that there was a much greater degree of overlap between the searchlight maps and regions that were more active for trials preceding safe choices (8,520 voxels of a possible 43,995 in the univariate mask) than between searchlight maps and regions that were more active for trials preceding risky choices (1,706 voxels of a possible 35,215 in the univariate mask) (Fig. 4), suggesting that there is greater local discriminative information in regions that respond more when avoiding a risk than when pursuing one. These findings indicate that the control networks identified in the searchlight analyses are more active before avoiding a risk, suggesting that cognitive control may be necessary to initiate a safe choice after a series of risky choices and that failing to engage the necessary cognitive control processes may lead to increased risk taking.

Discussion

The ability to avoid or pursue risk appropriately is critical to adaptive behavior, and it underlies a number of public health issues, such as drug use and reckless driving. In the present study, we examined whether it was possible to use data collected before the point of decision making to predict choice behavior in a naturalistic risk-taking task. We found that it is possible to predict with high accuracy whether the neural activation of an individual previously unseen by the classifier reflects safe or risky choice in subsequent risky decisions. Remarkably, this choice decoding was possible at comparable levels of accuracy even when each subject's data were reduced down to two values: one reflecting mean activation in risk-seeking regions and one reflecting mean activation in risk-avoidance regions.

Table 1. Regions where searchlight-based classification analysis discriminates between prepump and pre-cash out trials

Region	No. of voxels	MNI x	MNI y	MNI z	Peak classification, %
L parietal, L lateral occipital, L precentral and postcentral gyrus	4,617	−40.3	−33.1	47.9	72.2
R parietal, R lateral occipital	1,868	42.3	−45.2	47.4	67.6
R anterior insula, R lateral OFC	1,083	35.5	20.4	−0.5	70.4
R precentral gyrus	914	41.4	−0.1	38.2	69.9
mPFC	853	3.3	25.4	32.7	68.1
L anterior insula, L lateral OFC	778	−32.7	23.6	−2.7	69.9
Anterior cingulate, SMA	326	−4	−2	44	67.1

Montreal Neurological Institute (MNI) coordinates denote the 3D center of gravity of each cluster. L, left; mPFC, medial prefrontal cortex; R, right; SMA, supplementary motor area.

Helfinstein et al.

Table 2. Pearson correlations between searchlight classification map and NeuroSynth term-based reverse inference activation maps

Term	Correlation (r)
Control	0.1451
Working	0.1159
Numerical	0.1157
Letter	0.1081
Attention	0.1062
Correct	0.1060
Cue	0.0995
Preparatory	0.0970
Load	0.0959
Hand	0.0924

The 10 most highly correlated terms are listed.

Although our low-resolution analysis indicated that broadly distributed patterns of activation are sufficient to predict choice, our searchlight analyses indicated that local activation in a specific network of regions, namely, bilateral parietal and motor regions, the anterior cingulate cortex, bilateral insula, and bilateral lateral OFCs, show highly discriminable patterns of activation before risky choice and are also sufficient to predict risky choice. Together, these findings suggest that the neural computations underlying risky choices are reflected both in highly local signals and in coarse, broadly distributed patterns of activation, and they highlight the utility of highly engaging decision-making tasks that reliably correlate with real-world risk-taking behavior, such as the BART (7), in combination with approaches that allow us to answer targeted questions with these tasks.

The regions found to classify significantly between risky and safe choice in our searchlight analysis are highly consistent with regions found to be active in previous risk-taking studies, including the cingulate cortex, insula, parietal cortices, and striatum (1–6, 15–18). Of note, an analysis examining the overlap between regions identified in the searchlight analysis and regions used in the coarse-encoding analysis indicated that most of the identified searchlights were more active before a safe choice than before a risky choice. Nevertheless, the level of mean activation in regions more active before a risky choice alone was sufficient to predict choice significantly in the reduced-encoding analysis and provided significant predictive power beyond the mean activation in regions associated with risk avoidance. These results suggest that focusing analyses exclusively on regions with high local power causes us to overlook large networks that contain relevant, albeit coarsely encoded, information.

In addition to our use of a multivariate approach, our study differs from most previous risk research in its use of a paradigm that has strong external validity when it comes to real-world, health-relevant risk-taking behaviors. A major challenge when conducting risk-taking research is finding an appropriate way to operationalize risk, particularly given that economists define risk as the amount of variance in outcomes, whereas lay persons define risk as exposure to potential loss (7). Because the risky option in the BART has both greater variance in outcome and greater exposure to loss than the safe option, we managed to side step this issue by operationalizing risk in a way that encompasses both definitions. Most importantly, performance on the BART is highly correlated with real-world, public health-relevant risk taking, such as smoking, drug use, sexual risk taking, and unsafe driving behaviors (9, 11–14), suggesting that the cognitive processes engaged in the BART are comparable to those engaged in real-world risk taking. Many have shied away from the use of such complex tasks due to the possibility for confounds between different aspects of risk taking, but using the present analytical

approach allows us to take advantage of the validity of the BART while avoiding its potential complications.

Our searchlight analysis was designed to identify the regions, and thus the types of information encoding or cognitive processing, that best predict, and thus are most likely to influence, the outcome of a subsequent risky decision. The regions found to be active in searchlight analysis, and thus to provide the greatest local predictive information regarding risk taking, were shown in our reverse inference analysis to be regions typically involved in functions such as control, working memory, and attention. These regions were more active before safe choices in the univariate analysis, suggesting that successfully engaging executive processes may play a key role in initiating a safe choice. This role for executive control in initiating safe choices may be particularly pronounced due to the nature of the BART, where individuals typically make a series of increasingly risky choices before having to change their course of action to make a safe choice. However, many health-relevant risky decisions share this same structure, such as when deciding how many alcoholic beverages to drink before driving home or how much one can experiment with drugs or cigarettes before developing an addiction. Indeed, this characteristic of the BART may partially explain its strong external validity, and control structures may be critical for making the transition from engaging in mild risks to stopping when the risk becomes too great, both in the task and in real life.

In summary, the present study uses data from a large sample of subjects performing a naturalistic risk-taking task to determine whether blood-oxygen–level dependent data can be used successfully to predict choice behavior in subsequent risky decisions. We find that choice behavior can be predicted with a high level of accuracy and that the regions where activation patterns most successfully predict choice are areas known to be involved in executive functions, such as control, working memory, and attention, including the anterior cingulate cortex, bilateral insula, and parietal cortices. Choice can also be predicted with high accuracy using only estimates of mean activation across two large
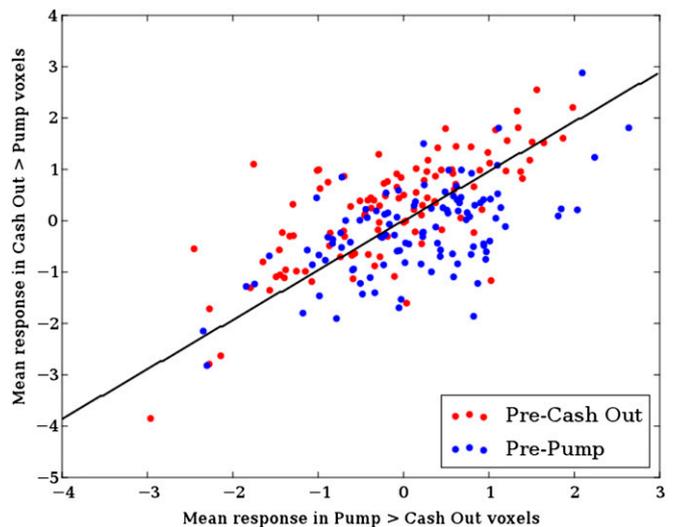
Fig. 3. Plot of sample values for the two parameters in the two-parameter classification. The x-axis values indicate mean Z-statistic values for the contrast between prepump and pre-cash out activity in regions where prepump activation is greater than pre-cash out activation ($t > 2.0$). The y-axis values indicate the same for regions where pre-cash out activation is greater than prepump activation. In each case, voxels were selected using independent training set data. Blue indicates prepump samples, and red indicates pre-cash out samples. The black line reflects the logistic regression classification boundary; this classifier was able to separate prepump and pre-cash out trials successfully (67% success rate).
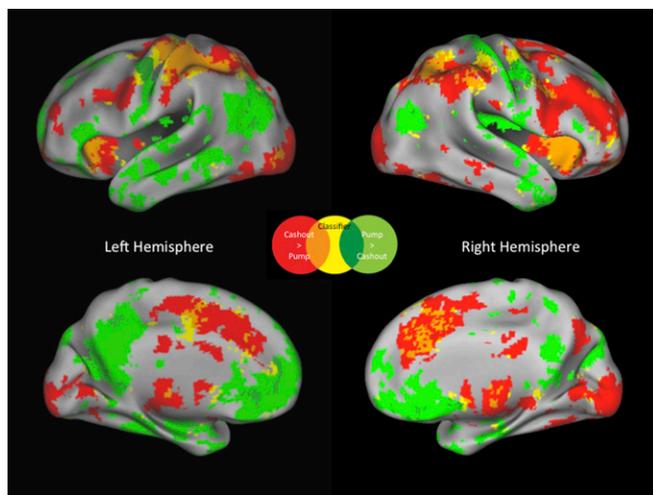
Fig. 4. Searchlights that could successfully discriminate between prepump and pre–cash out trials overlaid on a map of prepump vs. pre–cash-out activation thresholded at $t > 2.0$. More searchlights were located in regions where activation is higher on pre–cash out trials (8,520 voxels of 43,995 total voxels in the univariate mask) than on prepump trials (1,706 voxels of 35,215 total voxels in the univariate mask), suggesting that it is changes in regions related to risk aversion that most reliably predict whether a subject will make a risky or safe choice.

networks, indicating that information encoding risky decision-making behavior is represented at a very low spatial resolution, as well as in highly localized regions.

Overall, our findings suggest that clearly dissociable brain states precede selection of risky vs. safe choices in a risky decision-making situation, even when objective levels of risk and imminent motoric responses are identical. These patterns of activation differ in dissociable ways both in the overall balance of activation between systems active before risk seeking and in those active before risk avoidance, as well as in highly localized changes in activation, mostly in regions sensitive to risk avoidance.

Methods

Participants. All participants were recruited from the Los Angeles area as part of the Consortium for Neuropsychiatric Phenomics at the University of California, Los Angeles (UCLA) (www.phenomics.ucla.edu), with the goal of examining differences in response inhibition and working memory between healthy adults and patients diagnosed with bipolar disorder, schizophrenia, or adult attention deficit hyperactivity disorder (ADHD). Only data collected from healthy participants were included in the present analyses. All candidates had telephone screening, followed by additional in-person screening and informed consent. For both the healthy and patient groups, participants were men or women between the ages of 21 and 50 y who met the following selection criteria: a National Institutes of Health racial/ethnic category of either white/not Hispanic or Latino or Hispanic or Latino; a primary language of either English or Spanish; completion of at least 8 y of formal education; no significant medical illness; adequate cooperation to complete assessments; and visual acuity of 20/60 or better. Additional exclusion criteria for participants that took part in the imaging portion of the study included left-handedness, pregnancy, history of head injury with loss of consciousness or cognitive sequelae, or other contraindications to scanning (e.g., claustrophobia, metal in body, body too large to fit in scanner).

Participants in the healthy group were excluded if they had lifetime diagnoses of schizophrenia or another psychotic disorder; bipolar I or II disorder; or a current diagnosis of major depressive disorder, suicidality, anxiety disorder (obsessive-compulsive disorder, panic disorder, generalized anxiety disorder, or posttraumatic stress disorder), ADHD, or substance abuse/dependence.

All participants gave written informed consent according to the procedures approved by the Institutional Review Boards at UCLA and the Los Angeles County Department of Mental Health.

Procedure. A portion of the total sample took part in two separate functional MRI (fMRI) sessions, each of which included 1 h of behavioral testing and a 1-h scan on the same day (order counterbalanced across subjects). A total of 139 healthy adults completed at least a portion of the fMRI sessions; of those, 7 were excluded altogether [4 for missing a magnetic prepared rapid acquisition gradient-echo sequence (MPRAGE) scan, 2 for withdrawing from the study, and 1 for noncompliance with scan procedures] and 24 were excluded from the present BART analyses (13 for excessive motion, 4 for task performance that prohibited model fitting, and 7 for an unusable MPRAGE), leaving 108 participants included in the analyses.

Experimental Design. Participants performed a variant of the BART (9) while undergoing an fMRI scan (Fig. 1). Each round began with the presentation of either a green (experimental) or white (control) balloon. Experimental balloons constituted approximately two-thirds of the balloons presented. For experimental balloons, subjects were instructed to choose on each trial to either "pump" or cash out the balloon. Each successful pump would earn them five points, to be received when the subject cashed out the balloon. However, some pumps would cause the balloon to "explode," and if an explosion occurred before the subject cashed out, no points would be received for that balloon. Thus, each pump opportunity presented the subjects with a risky decision: They could select the safe option, a guaranteed receipt of the amount of points already earned if they cashed out, or the risky option, pumping, with a chance to receive an additional five points but also with a chance of losing the points already earned. Each balloon was programmed to explode on a particular pump number, drawn randomly from a uniform distribution between 1 and 12, although subjects were not informed of the probability structure of explosions. After subjects made a choice, a blank screen was presented for 1–2 s (average of 1.5 s), followed by presentation of the outcome. If the balloon exploded, an image of an exploded balloon appeared on the screen for 2 s. If not, a larger version of the balloon appeared and subjects had another opportunity to choose to either pump or cash out. After cash outs, subjects saw a screen with their accumulated earnings for that round for 2 s and the task then moved on to the next balloon. Subjects were instructed to make any button response to control balloons until they disappeared. Balloons were separated by a blank screen that was presented for 1–12 s (average of 4 s). Subjects had 9 min to complete the task at their own pace, seeing an average of 18.03 experimental balloons (range: 9–26).

Stimulus presentation and timing of all stimuli and response events were achieved using MATLAB (MathWorks) and the Psychtoolbox (www.psychtoolbox.org) (27) on an Apple PowerBook. For the experiment block administered in the scanner, each participant viewed the task through MRI-compatible goggles and responded with his or her right hand on an MRI-compatible button box in the scanner.

fMRI Data Acquisition. Data were collected at two scanning facilities at UCLA, each with a 3-T Siemens Trio MRI scanner, using a T2*-weighted echoplanar sequence [repetition time (TR) = 2.0 s, echo time (TE) = 30 ms, 64 × 64 matrix, field of view (FOV) = 192 mm, flip angle = 90°, slice thickness = 4.0 mm, 34 slices, oblique slice orientation, 267 volumes per run, one run]. A T2-weighted, matched-bandwidth, high-resolution structural scan with the same slice prescription as the functional images was also collected, and a T1-weighted, high-resolution, volumetric scan using a MPRAGE sequence (TR = 1.9 s, TE = 2.2 6 ms, 256 × 256 matrix, FOV = 250 mm, slice thickness = 1 mm, 176 slices) was used for anatomical registration.

Imaging data were processed and analyzed using the The Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB)'s Software Library (FSL). FMRIB's Model Correction Linear Image Registration Tool (MCFLIRT) was used for motion correction. Following motion correction, Multivariate Exploratory Linear Optimized Decomposition into Independent Components (MELODIC) was used to identify artifactual components in the data, which were modeled as nuisance regressors. Brain Extraction Tool (BET) was used to extract the brain from the skull. The data were filtered using a nonlinear high-pass filter with a 100-s cutoff. Data were smoothed using a 5-mm FWHM isotropic Gaussian filter.

fMRI Data Analysis. To conduct the classification analysis, patterns of activation on two types of trials were compared: those that immediately preceded a cash out trial and those that immediately preceded a pump trial (Fig. 1). For each subject, all pump trials that immediately preceded a cash out trial were identified. Each of these trials was then matched with a pump trial that immediately preceded a pump trial and had the same pump opportunity as the cash out trial (i.e., the number of times the subject had already pumped the balloon when this trial occurred was the same for the prepump and pre-

cash out trials). The prepump trials were randomly chosen from the pools of appropriate trials using Python's random.sample function. All prepump and pre-cash out trials that could not be matched in this way were not included in the prepump and pre-cash out regressors. Thus, the distribution of risk (i.e., amount of points that could be lost and likelihood of explosion) for the pre-pump and pre-cash out regressors was identical. The selected prepump and pre-cash out trials were then included as regressors in the linear regression model. To control for differences in reaction time between prepump trials and pre-cash out trials, we used a variable epoch regressor, where the duration of each regressor was the length of the subject's reaction time on that trial (24). In addition to these regressors, all pump and cash out trials were modeled with two regressors: one with a fixed amplitude and a fixed duration equivalent to the average reaction time (RT) for all pump opportunities and one with a fixed amplitude and duration equivalent to the actual RT on that pump opportunity. All control balloon pump trials were modeled with two analogous regressors. Six motion parameters and their temporal derivatives were also included as regressors of no interest. Beta-weights from each subject for the prepump and pre-cash out regressors were then used as inputs for the machine-learning algorithm. Whole-brain and searchlight (25) classifications were conducted using a linear support vector machine classifier in PyMVPA (28). Classification was cross-subject, with subjects divided into six folds, with classification of the subjects in each fold predicted based on classification of subjects in the other five folds. For whole-brain analysis, this process was repeated 50 times with subjects randomly assigned to a fold each time. Searchlight results were thresholded at a classification rate of 60%, and a random permutation procedure was then used to determine the appropriate cluster size for correction at $P = 0.05$. Searchlight maps were then correlated with term-based maps in NeuroSynth (26) to determine the cognitive processes most strongly associated with searchlight regions.

After searchlight analyses, we conducted a follow-up analysis examining classification using coarse measures of activation in regions where activation differed substantially between pre-cash out and prepump conditions. To do this, we used only training samples (five-sixths of the total data) to again determine voxels where the difference in activation between cash out and pump trials was $t > 2.0$ (prepump > pre-cash out) or $t < -2.0$ (pre-cash out > prepump). For each training sample, we then calculated the mean activation across the $t > 2.0$ voxels and the mean activation across the $t < -2.0$ voxels, reducing the information in each training sample to two values. We fed these two values into a logistic classifier, which we then used to classify test samples based on the mean activation from the same two sets of voxels. We repeated this process 50 times with subjects randomly assigned to folds on each iteration and then determined statistical significance via random permutation.

We then examined whether searchlights that enabled significant classification were located primarily in regions that showed greater activation before cash out or in regions that showed greater activation before pumps. To do this, we used the same prepump and pre-cash out regressors used in the classification analysis and calculated univariate contrast maps, with per-voxel thresholds of $t > 2.0$ (prepump > pre-cash out) and $t < -2.0$ (pre-cash out > prepump), this time using data from all subjects. We then calculated the number of voxels significant in the searchlight analyses that overlapped with each map.

1. Mohr PNC, Biele G, Heekeren HR (2010) Neural processing of risk. *J Neurosci* 30(19): 6613–6619.
2. Krain AL, Wilson AM, Arbuckle R, Castellanos FX, Milham MP (2006) Distinct neural mechanisms of risk and ambiguity: A meta-analysis of decision-making. *Neuroimage* 32(1):477–484.
3. Preuschoff K, Bossaerts P, Quartz SR (2006) Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51(3):381–390.
4. Rudorf S, Preuschoff K, Weber B (2012) Neural correlates of anticipation risk reflect risk preferences. *J Neurosci* 32(47):16683–16692.
5. Rao H, Korczykowski M, Pluta J, Hoang A, Detre JA (2008) Neural correlates of voluntary and involuntary risk taking in the human brain: An fMRI Study of the Balloon Analog Risk Task (BART). *Neuroimage* 42(2):902–910.
6. Schonberg T, et al. (2012) Decreasing ventromedial prefrontal cortex activity during sequential risk-taking: An FMRI investigation of the Balloon Analog Risk Task. *Front Neurosci* 6:80.
7. Schonberg T, Fox CR, Poldrack RA (2011) Mind the gap: Bridging economic and naturalistic risk-taking with cognitive neuroscience. *Trends Cogn Sci* 15(1):11–19.
8. Fox CR, Tannenbaum D (2011) The elusive search for stable risk preferences. *Front Psychol* 2:298.
9. Lejuez CW, et al. (2002) Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task (BART). *J Exp Psychol Appl* 8(2):75–84.
10. Weber EU, Johnson EJ (2008) Decisions under uncertainty: Psychological, economic, and neuroeconomic explanations for risk preference. *Neuroeconomics: Decision Making and the Brain*, eds Glimcher PW, Camerer CF, Fehr E, Poldrack RA (Elsevier, London), pp 127–144.
11. Aklin WM, Lejuez CW, Zvolensky MJ, Kahler CW, Gwadz M (2005) Evaluation of behavioral measures of risk taking propensity with inner city adolescents. *Behav Res Ther* 43(2):215–228.
12. Hopko DR, et al. (2006) Construct validity of the Balloon Analogue Risk Task (BART): Relationship with MDMA use by inner-city drug users in residential treatment. *J Psychopathol Behav Assess* 28(2):95–101.
13. Lejuez CW, Aklin WM, Zvolensky MJ, Pedulla CM (2003) Evaluation of the Balloon Analogue Risk Task (BART) as a predictor of adolescent real-world risk-taking behaviours. *J Adolesc* 26(4):475–479.
14. Lejuez CW, Simmons BL, Aklin WM, Daughters SB, Dvir S (2004) Risk-taking propensity and risky sexual behavior of individuals in residential substance use treatment. *Addict Behav* 29(8):1643–1647.
15. Engelmann JB, Tamir D (2009) Individual differences in risk preference predict neural responses during financial decision-making. *Brain Res* 1290:28–51.
16. Preuschoff K, Quartz SR, Bossaerts P (2008) Human insula activation reflects risk prediction errors as well as risk. *J Neurosci* 28(11):2745–2752.
17. Weber BJ, Huettel SA (2008) The neural substrates of probabilistic and intertemporal decision making. *Brain Res* 1234:104–115.
18. Symmonds M, Wright ND, Bach DR, Dolan RJ (2011) Deconstructing risk: Separable encoding of variance and skewness in the brain. *Neuroimage* 58(4):1139–1149.
19. Paulus MP, Rogalsky C, Simmons A, Feinstein JS, Stein MB (2003) Increased activation in the right insula during risk-taking decision making is related to harm avoidance and neuroticism. *Neuroimage* 19(4):1439–1448.
20. Matthews SC, Simmons AN, Lane SD, Paulus MP (2004) Selective activation of the nucleus accumbens during risk-taking decision making. *Neuroreport* 15(13): 2123–2127.
21. Kuhnen CM, Knutson B (2005) The neural basis of financial risk taking. *Neuron* 47(5): 763–770.
22. Lee TMC, Leung AWS, Fox PT, Gao J-H, Chan CCH (2008) Age-related differences in neural activities during risk taking as revealed by functional MRI. *Soc Cogn Affect Neurosci* 3(1):7–15.
23. Xue G, et al. (2009) Functional dissociations of risk and reward processing in the medial prefrontal cortex. *Cereb Cortex* 19(5):1019–1027.
24. Grinband J, Wager TD, Lindquist M, Ferrera VP, Hirsch J (2008) Detection of time-varying signals in event-related fMRI designs. *Neuroimage* 43(3):509–520.
25. Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci USA* 103(10):3863–3868.
26. Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods* 8(8): 665–670.
27. Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10(4):433–436.
28. Hanke M, et al. (2009) PyMVPA: A python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics* 7(1):37–53.
29. Van Essen DC (2005) A Population-Average, Landmark- and Surface-based (PALS) atlas of human cerebral cortex. *Neuroimage* 28(3):635–662.