

# Supporting Information

Vilares et al. 10.1073/pnas.1619385114

## SI Materials and Methods

### Experimental Details.

**Experimental paradigm.** Participants were told a cover story about carrying “valuable content” (such as documents, microchip processors, etc.), here referred to as “contraband,” through a checkpoint (Fig. S1). Note that, although the instructions did not use the term contraband so as not to discourage participants that were averse to illegal behavior, we use the term here for convenience. In each trial, they were shown between one and five suitcases, only one of which actually contained contraband, and were asked whether they were willing to carry a suitcase randomly chosen from the group (Fig. S1A, *Left*). Hence, the number of suitcases shown represented the risk of carrying the target suitcase with contraband (Contraband Risk): if only one suitcase was presented, then the participants knew with certainty that the suitcase had contraband in it (knowing situation,  $P_{\text{contr}} = 1$ ), whereas if more than one suitcase was presented they were not sure whether the suitcase they were assigned contained contraband, but were aware of the risk (reckless situation, with  $P_{\text{contr}} = 0.5, 0.33, 0.25$ , or  $0.2$  of having contraband in the suitcase). Participants also had different probabilities of being caught (Search Risk), with the probability of being searched at the checkpoint ranging from  $P_{\text{search}} = 0$  to  $0.8$  (symbolized by 10 tunnels, in which a proportion of them could be occupied by a “guard”; Fig. S1A, *Right*). One-half of the participants ( $n = 20$ ) saw the probability of carrying a suitcase with contraband after already being shown the search risk (Search-First group), whereas the other half started by seeing the suitcases before being shown the search risk (Contraband-First group).

The motivation to have a Contraband-First condition and a Search-First condition was to control for the order of presentation of the information. The motivation to not only change the contraband risk but to change the search risk as well was twofold. First, it enabled us to provide a more realistic setup—one that involved (as in real life) differing potential search risks. Second, it enabled us to vary search risks together with contraband risks, in a way that allowed us to disambiguate the effects of Contraband Risk, Expected Value, and Variance in Reward. Without variation in search risk, these would be perfectly correlated, and therefore indistinguishable.

Task payouts are shown in Fig. S1B. At the beginning of each trial a participant was endowed with \$6,000. After seeing the two types of risk information (Contraband Risk and Search Risk), participants decided whether or not to carry a suitcase. To carry a suitcase, participants had to pay \$500. If they decided not to carry, they had to pay \$1,500, leaving them with a trial total of \$4,500. If no decision was made (participants did not respond during the allocated time), participants lost \$2,500, leaving them with a trial total of \$3,500. These three costs associated with each trial (the cost to carry, the cost to reject, and the cost of inaction) were considered a “life tax” and were included (informed by pilot studies) to simulate motivations to act expressed by individuals engaging in the real-world analog of our experiment. If participants chose to carry a suitcase, if it contained contraband (“target suitcase”) and if they were not searched, they received \$2,500 extra, getting a trial total of \$8,000 (the maximum possible payoff). However, if participants were searched while carrying the target suitcase containing contraband, they lost \$3,500, leaving them with a trial total of \$2,000 (the minimum possible payoff). If they chose to carry a suitcase and it did not contain contraband (“dummy suitcase”), they received no extra money, leaving them with a trial total of \$5,500, regardless of whether or

not they were searched. Participants were not shown the results of individual trials. Hence, because no feedback was provided at the end of each trial, all trials were independent from one another. At the end of each trial, a computer simulated the outcome for that trial. At the end of the experiment, the computer randomly picked the payout of one trial and participants earned 1% of their trial total in US dollars. Hence, each participant received \$20 to \$80 at the end of the experiment, the exact value depending on the trial randomly chosen by the computer and the choices and outcome of that trial.

**Display/stimulus.** The sequence of each trial was as follows (Fig. S1C): First, the contraband risk (Contraband-First group) or the search risk (Search-First group) was shown on the screen, for 2 s. Afterward, there was a blank screen for about 3 s (duration jittered between 2.5 and 3.5 s), and the search risk or contraband risk (respectively) was displayed, also for 2 s. Then another blank screen was shown (duration: average 3 s, jittered between 2.5 and 3.5 s), and a screen appeared indicating that the program was selecting which suitcase would be carried (duration: 0.75 s), after which another screen was shown asking the subject to choose to carry that suitcase or not. Participants had free time to make their decision, and they expressed their decision by pressing a button (representing either “yes” or “no”). If no button was pressed in the following 5 s, it was recorded as if no decision was made. After the decision was inputted (or the 5 s passed), the choice made by the participant was highlighted on the screen, and the computer calculated the results of that trial/round (not shown to the participant). Finally, about 0.75 s after the choice was submitted, participants were shown a screen with a fixation cross, which lasted for about 3 s (between 2.5 and 3.5 s, jittered), and then a new trial started.

Participants performed a total of 125 trials in the experiment (5 types of Search Risk  $\times$  5 types of Contraband Risk  $\times$  5 repetitions of each trial type). The exact sequence of trials shown was chosen randomly and differed among participants. The sequence of each trial can be seen in Fig. S1C.

### Behavioral Data Analyses.

**Statistical tests.** To determine whether there were significant differences in the choice to carry contraband based on the risk of having contraband, the risk of being searched, and the order of risk information received, a  $5 \times 5 \times 2$  mixed-model ANOVA was conducted with Contraband Risk and Search Risk serving as within-group factors containing five levels each and the order of information seen serving as the two-level between-group factor (Contraband-First or Search-First). This analysis was done using the statistical software SPSS.

**Logistic regression.** The participants’ behavior was analyzed using logistic regression as implemented by the function `glm` in R (Team 2014). The dependent variable was a participant’s response: Accept or Reject carrying the suitcase. The dependent variables were Contraband Risk, Search Risk, and the interaction Contraband Risk\*Search Risk. The condition (Contraband-First or Search-First) was also entered as a categorical variable, with Contraband-First decisions coded as 0, and Search-First decisions as 1. The terms in the model were thus as follows: a constant, a dummy variable for Condition, Contraband Risk, Contraband Risk\*Condition, Search Risk, Search Risk\*Condition, Contraband Risk\*Search Risk, and Contraband Risk\*Search Risk\*Condition.

### fMRI Analysis.

**Scans and preprocessing.** Anatomical and functional images were acquired on a Siemens 3-T Trio scanner at Virginia Tech Carilion



level-dependent response for that experiment/participant) and also six nuisance regressors, which corresponded to participant-specific head-movement parameters. Contrast images, derived from a pairwise contrast between each event type and an implicit baseline (equivalent to the beta image produced for that event), were then calculated for each event and used as baseline data for the elastic-net (EN) regression.

### EN Regression.

**Data used to estimate the model.** For each participant, we extracted two contrast/beta images (obtained by the GLM procedure outlined above): one belonging to the event type knowing (when only one suitcase was shown, and so participants knew that they had the target suitcase with contraband, i.e.,  $P = 1$  of having the target), and one belonging to a reckless situation (e.g., when five suitcases were presented, i.e., probability of having the target suitcase with contraband is  $P = 0.2$ ). The brain data corresponding to each contrast condition are then reshaped into a vector ( $1 \times$  number of voxels) and entered as a row in the data matrix. After doing this for each participant, the final data matrix that will be used for EN regression has  $2n$  rows, with  $n$  representing the number of participants and 2 representing the number of different events we are trying to predict; and  $p$  voxels, corresponding to the number of voxels in the participant's contrast image. In our task, the matrices had 40 rows (20 participants in each group  $\times$  2 events) and 65,280 columns (each one corresponding to one voxel in the brain). We then took out the voxels that did not fall inside the brain, leading to a final number of voxels/columns of  $\sim 21,000$ . Thus, for the EN model, we had 40 observations (2 observations per participant) and 21,000 predicting variables (i.e., features). Each observation was associated with one particular label (knowing or reckless), which we then tried to predict.

We chose to model knowledge vs. recklessness at the time the contraband risk is revealed to have a “cleaner” knowledge vs. reckless brain state, and so maximize our chances of observing these two brain states, should they exist.

**The EN regression.** The EN regression (21) is a form of regularized linear regression that tries to minimize as follows:

$$\min_{\beta_0, \beta} \frac{1}{N} \sum_{i=1}^N l(y_i, \beta_0 + \beta^T x_i) + \lambda \left[ (1 - \alpha) \|\beta\|_2^2 / 2 + \alpha \|\beta\|_1 \right],$$

where  $y_i$  is the vector we are trying to predict (in our case, composed by two labels, knowing and reckless),  $x_i$  are the predictor variables (in our model, the voxels),  $N$  corresponds to the number of observations,  $l(\dots)$  is the loss function associated with the type of data we are using (e.g., Gaussian, binomial);  $\beta$  corresponds to the coefficients of the model that we are trying to estimate (being  $\beta_0$  the intercept),  $\lambda$  (lambda) is a regularization parameter that controls the strength of the regularization (for high values of  $\lambda$ , all of the coefficients will tend to zero; for  $\lambda = 0$ , the EN regression becomes an ordinary least-squares regression), and  $\alpha$  (alpha) is the EN mixing parameter, which varies from 0 to 1 and indicates the relative quantities of L2 norm penalized regression (ridge regression, which corresponds to an  $\alpha = 0$ ) and L1 norm penalized regression (lasso regularized regression, which corresponds to an  $\alpha = 1$ ). Having some kind of regularization of the  $\beta$  coefficients/predictors is very important to minimize overfitting in the case in which the number of predictors ( $p$ ) is much greater than the number of participants ( $n$ ) (i.e.,  $p \gg n$ ) (14), as is the case with fMRI data (if we assume that each voxel is a variable). Ridge regression ( $\alpha = 0$ ) does not make variable selection, so although it shrinks the  $\beta$  coefficients, it keeps all of them, hence it becomes harder to disentangle which predictor variables are important. Lasso regression, on the other hand, does make variable selection but it does not allow for

several correlated coefficients to remain in the final model, even if these coefficients are also important for the model. Furthermore, it never retains more variables than the number of observations. The EN penalized regression, by mixing these two types of penalization, is able to make variable selection while at the same time allowing for clustering of potential relevant variables. **Parameter selection and model fitting.** To fit the EN penalized logistic regression, we used the glmnet package for Matlab ([web.stanford.edu/~hastie/glmnet\\_matlab/](http://web.stanford.edu/~hastie/glmnet_matlab/)), which also included a function to do cross-validated model fitting (cvglmnet.m) and a function to make predictions (glmnetPredict.m). To select the two tuning parameters for the EN, the mixing parameter ( $\alpha$ ) and the overall complexity parameter ( $\lambda$ ), we first estimated  $\alpha$  and, after  $\alpha$  was fixed,  $\lambda$  was selected (14) (Fig. S2). To estimate  $\alpha$ , similar to the procedure used by Ahn et al. (14), we conducted a grid search over different values of  $\alpha$  (from 0 to 1, in small increments).

- i) For each potential  $\alpha$  value, we did the following procedure:
  - a) Fivefold cross-validated EN fitting, using the cvglmnet.m function. What this function does, step by step, is as follows:
    - i) At the particular value of  $\alpha$  chosen, it first fits the EN model paths to get the  $\lambda$  sequence.
    - ii) It then divides the data randomly into five groups, or folds.
    - iii) For each value of  $\lambda$ , it fits a model (in this case, a logistic model) using penalized maximum likelihood, that is, it fits an EN model, using only 80% of the data (four of the five folds previously made).
    - iv) It then tests the model on the left-out fold, and computes the corresponding minimum binomial deviance. It does that for each  $\lambda$ .
    - v) Steps i, ii, iii, and iv are then repeated five times (one for each fold), and the average binomial deviance over the five repetitions is computed for each  $\lambda$ .
  - b) For each of these runs (performed by cvglmnet), we then saved the minimum average binomial distance obtained with the model.
  - c) For each  $\alpha$ , we repeated 40 times steps a and b, and calculated the mean of the minimum average binomial deviance over the 40 repetitions.
- ii) We then chose the  $\alpha$  that minimized this value.
- iii) After choosing  $\alpha$ , we proceeded to estimate  $\lambda$ , and the  $\beta$  coefficients of the model. We repeated the following procedure 40 times:
  - a) We ran a fivefold cross-validated EN regression (identical to step 1a), using the  $\alpha$  obtained from steps i and ii.
  - b) From it, we obtained the  $\lambda$  value that minimizes the binomial deviance ( $\lambda_{\min}$ ).
  - c) We also obtained the indices associated with the division of the data in five folds.
  - d) Then, for  $\lambda_{\min}$ , we:
    - i) Refitted an EN regression using only four folds of the data (using the division made by the cvglmnet function) and extracted the corresponding  $\beta$  coefficients for all regressors (voxels);
    - ii) Recorded, for each regressor (voxel), if it “survived” in the current run, that is, if its  $\beta$  value is not zero.
    - iii) We then used the resulting EN model and tested it on the “left-out” fold, obtaining the “fitted probabilities” associated with each observation in the left-out fold, and also the corresponding most likely classification (i.e., if the observation most likely came from a knowing or reckless situation). From



these, we computed both the correct classification rate (CCR; how often did it accurately guess the situation the observation belonged to) and the area under the curve (AUC) of the receiver-operating characteristic curve. Note that the CCR and the AUC computed this way represent the out-of-sample, cross-validated performance of the model, as the observations in the left-out fold were not used to calculate the EN model, and are suggestive of how well the model would generalize.

- iv) Repeated steps *i* through *iii* five times, each time leaving out a different fold.
- iv) After doing steps *a* through *d* 40 times, we then calculated a matrix with the “signed survival rate” for each voxel. For that, we calculated the “survival rate” of each voxel, meaning how many times that voxel “survived” on the total of the 200 runs done (40 iterations  $\times$  5 folds), and we multiplied this survival rate with the sign of the mean  $\beta$ -coefficient values obtained for that voxel.
- v) We then projected the signed survival rate of each voxel back into the brain. An average positive survival rate indicates voxels more predictive of being in the reckless state (i.e., the estimated coefficient values associated with those voxels were on average positive), whereas a negative survival rate indicates voxels more predictive of being in a knowing situation (which, in our model, was the “baseline” label). Higher survival rates indicate that the voxel was used frequently to distinguish the knowing and reckless situations.
- vi) To calculate the overall performance of the model, we made the average of the 200 AUCs and CCRs obtained during step 3 (40 iterations  $\times$  5 folds per iteration).

**Leave-one-subject-out cross-validation.** To obtain single-subject accuracy, we performed the same general steps as outlined above but using data from only  $n - 1$  participants each time. Specifically, in each iteration, we leave the data of one participant out and do a fivefold CV to obtain both  $\alpha$  and  $\lambda$  (using the same general steps as outlined above). We then fit an EN model using the data of the  $n - 1$  participants and the parameters ( $\alpha$  and  $\lambda$ ) obtained through fivefold CV, extract the corresponding betas, and test the data on the left-out participant. This procedure is then repeated  $n$  times (equal to the number of participants).

**Double-cross-validation procedure.** To obtain a very stringent measure of out-of-sample performance, we did a “double-cross-validated” procedure. Namely, we divided the participants in half, randomly, and built an EN classifier using only one-half of the data (doing fivefold cross-validation of the parameters within that data, following the same general steps outlined above), and then tested the resulting model on the untouched other half of the data. This other half of the data then serves as a completely independent dataset, as it was not used at all to train the EN classifier. This process is then repeated for the other half and the results averaged. For the outer fivefold inner-fivefold double cross-validation, the procedure is exactly the same, just that the participant’s data are divided in five (instead of two) groups, and the EN classifier fitted on four-fifths of the data (using fivefold CV within it to fit it) and tested on the remaining one.

**Permutation test.** If there is no real information present in the data, a binary classifier (as in the example here, given that we are trying to classify two different labels, knowing and reckless) should give, on average, a correct classification rate of 50% and an AUC of 0.5. Hence, anything above that should indicate that the model is performing better than chance. However, these values only hold for infinite sample sizes, and for small sample sizes the values can be higher (or lower) by chance (23). To assess the “significance” of the results we obtained in the models, we ran a permutation test, in which, for each permutation run, we followed the exact

same procedure as outlined above, but in which we shuffled the labels corresponding to each observation (i.e., in our experiment, the labels that said if the data belonged to a knowing or a reckless situation). We iterated this procedure 200 times. Then, to obtain the “*P* value” for a particular AUC/CCR value, we calculated the proportion of iterations that had an AUC/CCR that high or higher. The results obtained with the CCR were identical to the ones presented in Combrisson and Jerbi (23) and very close to the theoretical values (based on a binomial cumulative distribution function, for number of observations = 40 and 2 different groups to be distinguished), with “statistical significance” (i.e., a *P* value of 0.05) achieved on average around AUC = 0.70/CCR = 60%. Note: for the permutation test, we fixed  $\alpha = 0.5$  (due to time constraints). Redoing all of the analysis reported in the text using  $\alpha = 0.5$  leads to the same basic results.

## SI Results and Discussion

**Behavioral Analysis—Logistic Regression.** The logistic regression over participants’ behavior revealed that, even when the monetary gain/loss potential was equal, individuals who processed the likelihood of being searched by guards first (Search-First condition), followed by the likelihood that a suitcase contained contraband, were less likely to carry a contraband suitcase compared with those who processed the same information in the opposite order (Contraband-First condition; Table S1).

More precisely, logit regression analysis showed (*i*) a significant main effect of Search Risk; (*ii*) a significant Contraband Risk by Search Risk interaction; and (*iii*) a significant Condition by Contraband Risk by Search Risk interaction: The effect of Contraband Risk became greater as Search Risk increased, and this effect was bigger when the Search Risk was seen first. This order-dependent behavioral effect can be seen as a temporally extended framing effect. It is well known that human decision-making can be influenced by the manner in which options are presented (27, 28). Our results suggest that this is true not only for decisions involving multiple options but also for differing presentations of information related to a single decision. In the context of the current task, it is plausible that the likelihood of being searched is a more aversive signal compared with the likelihood that a case being carried might contain contraband. Our results indicate that this signal, when processed before further information arrives, increases the impact of Contraband Risk, making knowing (that is, Contraband Risk = 1) even more salient.

**Control Analyses.** To further confirm that the high performance of the EN classifier in the Search-First condition is not just driven by differences in visual information, we reran the obtained EN classifier (distinguishing one vs. five suitcases) but taking away all of the surviving voxels that were part of the occipital/visual cortex. The resulting EN classifier maintained its high predicting ability, having an out-of-sample average AUC value of 0.834 ( $P_{\text{perm}} = 0.005$ ) and an average CCR of 70.6% ( $P_{\text{perm}} = 0.005$ ), suggesting once more that it is not the visual information driving this high performance of the classifier.

**Single-Subject Precision.** To obtain a measure of single-subject precision, we fitted an EN classifier using a leave-one-subject-out procedure on top of the fivefold cross-validation (see *SI Materials and Methods, EN Regression, Leave-One-Subject-Out Cross-Validation* for details). We found that, for the Search-First condition, the EN classifier was able to predict with high accuracy whether the brain data corresponded to a knowing (Contraband Risk:  $P_{\text{contr}} = 1$ ) or a reckless ( $P_{\text{contr}} = 0.2$ ) situation. The EN classifier had an out-of-sample mean AUC of 0.944 ( $P_{\text{perm}} = 0$ ) and a mean correct classification rate of 81.8% ( $P_{\text{perm}} = 0$ ). For the Contraband-First condition, the EN classifier had an out-of-sample mean AUC of 0.499 ( $P_{\text{perm}} = 0.33$ ) and a mean correct classification rate of 50% ( $P_{\text{perm}} = 0.1$ ). Thus, in our Search-First condition, we were able

to obtain a high single-subject precision in distinguishing a knowing from a reckless scenario.

**Half-Split/Double-Cross-validation.** Within each group, we also split the participants in half, randomly, and built an EN classifier using only one-half of the data (and doing fivefold cross-validation of the parameters within that data). We then tested the resulting model on the untouched other half of the data. Using this extra, more stringent, analysis, we still observe a higher-than-chance prediction accuracy in the Search-First condition. Specifically, the EN classifier achieved an out-of-sample mean AUC of 0.765 ( $P_{\text{perm}} = 0$ ) and a mean correct classification rate of 73.9% ( $P_{\text{perm}} = 0$ ). For the Contraband-First condition, the EN classifier had an out-of-sample mean AUC of 0.503 ( $P_{\text{perm}} = 0.46$ ) and a mean correct classification rate of 50.5% ( $P_{\text{perm}} = 0.15$ ). Similarly, when we split the data into five groups, train the classifier on four-fifths of the data (using fivefold CV within this group), and then test on the left-out data, we also obtain good prediction accuracies. This approach yielded a mean AUC of 0.803 ( $P_{\text{perm}} = 0$ ) and a mean CCR = 73.3% ( $P_{\text{perm}} = 0$ ). Hence, the higher-than-chance prediction accuracies observed in the Search-First condition when classifying knowing vs. reckless are maintained even when using very stringent analyses.

**Additional Analyses.** The lack of predictive power for the EN in the Contraband-First condition is somewhat surprising, given that the only thing that changed between conditions was the order of presentation of information. A trivial explanation could be that the functional imaging data of one or more participants' in the Contraband-First condition is corrupted. If that were the case, then it would not be possible to obtain any good predictive models with this dataset. However, when predicting participants' decision to carry the suitcase (see below), the EN model performed with very high accuracies (AUC > 0.9) both in the Search-First group but also in the Contraband-First group. This indicates that there is some other reason for the low performance in the Contraband-First condition. Our behavioral analysis revealed that, although the content and the level of risk associated with a single decision were identical, the order in which the information was received significantly altered choice behavior (see *Results*, *Behavioral Results*, and also below). Thus, in our task, the order in which participants received information about contraband and search risk affected both their behavior and the corresponding imaging data.

**Knowing vs. Recklessness (Degree of Recklessness Not Specified).** We were also interested in understanding whether knowing could be broadly distinguished from reckless, even if no additional information about the degree of recklessness were given. For that, we built an EN classifier in which we contrasted knowing states with a general reckless state (which included all reckless states), not specifying the degree of recklessness (see *SI Materials and Methods* for details). We found that, for the Search-First condition, we were still able to predict with high accuracy whether the brain data corresponded to a knowing or reckless scenario, obtaining an average AUC value of 0.872 and a CCR of 75.9% ( $P_{\text{perm}} = 0.005$  and  $P_{\text{perm}} = 0$ , respectively). Thus, knowing seemed to be broadly distinguishable from reckless, even in the absence of information about the degree of recklessness.

**Knowing vs. Recklessness (No Search Risk).** To understand better whether the identified brain pattern was specifically associated with the extent of knowledge the participant had about the existence of contraband in the suitcase or whether the brain state only exists when the participant knows that there is a risk of getting searched, we reran the analysis but using only the knowledge and recklessness trials in which search risk was 0 (no guards on the tunnels). Hence there was no probability of getting

caught while carrying the contraband. Behaviorally, we can also see that they were aware the search risk was 0, as participants almost always decided to carry the suitcase when there was no risk of being searched (Fig. 1). We find that, for the Search-First condition, the EN model using only the no-search-risk trials did a good job in classifying a knowledge vs. recklessness scenario, giving a mean AUC = 0.832 and a mean CCR = 70.9% ( $P_{\text{perm}} = 0$  for both). Moreover, the brain areas obtained related to a knowing scenario were identical to the ones obtained in the full model [mPFC, cingulate cortex, insula, temporoparietal junction (TPJ); Fig. S3A]. Hence, the brain pattern we identified associated with the state of knowledge appears even if there is no threat.

#### Knowing vs. Recklessness or Just Difference in Number of Suitcases?

Finally, we wanted to know if the high accuracy results obtained with the classifier on the Search-First condition were mainly due to it being able to distinguish any linear increase in the number of suitcases being presented, and not to anything specific about the knowing/reckless distinction. If that is the case, then the classifier should perform similarly in distinguishing, say, two vs. five suitcases or distinguishing one vs. four suitcases. However, the EN classifier built to disentangle between two suitcases vs. five suitcases being presented (using the same procedure as before) did not perform better than chance, having an out-of-sample average AUC value of 0.243 and an average CCR of 30.6% ( $P_{\text{perm}} = 1$  for both). Contrast these values with the average values obtained at distinguishing between one vs. four suitcases (AUC = 0.82 and CCR = 75.7%,  $P_{\text{perm}} = 0$  for both). This once more indicates that the obtained high accuracy results (in the Search-First condition) are not simply due to a visual increase in the number of suitcases being presented, and suggests that there may be something special about the knowing/reckless distinction.

#### Brain Areas Specifically Associated with Contraband Risk, Expected Value, and Variance in Reward.

To try to understand whether the brain differential activations we observed between knowing and reckless were related to differences in Contraband Risk (knowing,  $P = 1$  of existence of contraband in the suitcase vs. reckless,  $0 < P < 1$ , aware of a possibility but not certainty of the existence of contraband) or whether they were just related to differences in Expected Value or Variance in Reward ["risk" as defined by the neuroeconomics literature (7–9)], we reran the same analyses but extracting out the effects associated with either of these factors (by modeling them separately at the first-level GLM model and not including them in the input data for the EN regression). We find that we still have a higher-than-chance accuracy in predicting a knowing vs. reckless scenario, with the EN regression having an out-of-sample CCR = 71.4% and an AUC = 0.791 ( $P_{\text{perm}} = 0.005$ ; Fig. S3B) when extracting out the potential effects associated with Variance in Reward, and an out-of-sample CCR = 71.7% and an AUC = 0.792 ( $P_{\text{perm}} = 0.005$ ; Fig. S3C) when excluding the effects associated with expected reward. Furthermore, if we do a simple GLM modeling independently Contraband Risk, Variance in Reward, and Expected Value, we see that the brain pattern we observed to be related to the knowing/reckless distinction continues to be specifically associated with Contraband Risk (probability of carrying the suitcase with contraband; see also Figs. S4 and S5). The EN model also performs well after taking out potential effects associated with the probability of being "caught," that is, of being searched while carrying the suitcase with contraband (CCR = 71.8%, AUC = 0.792,  $P_{\text{perm}} = 0.005$ ; Fig. S6), or the probability of getting the highest reward, that is, carrying the target suitcase and not being searched (CCR = 71.7%, AUC = 0.792,  $P_{\text{perm}} = 0.005$ ; Fig. S7). Together, this indicates that the general brain pattern we see associated with the knowledge/recklessness distinction seem to be specifically related to the probability of carrying the suitcase that has contraband (Contraband Risk) and cannot be explained only

by differences in expected value, variance in reward, fear of being searched, or expectation of highest reward.

**Simple GLM Results—Search Risk and Contraband Risk.** To understand which brain areas are specifically involved in signaling search risk and contraband risk, we added separate regressors for Contraband Risk and Search Risk in a traditional GLM model and analyzed the areas that were parametrically correlated with them.

The areas we found positively correlated with Search Risk were mainly in the visual cortex, namely the calcarine sulcus [ $P < 0.05$ , family-wise error (FWE) corrected]. Areas more active with decreasing Search Risk include the bilateral TPJ, dorsolateral prefrontal cortex (DLPFC), and middle temporal gyrus ( $P < 0.05$ , cluster size FWE corrected). Interestingly, the bilateral TPJ and middle temporal gyrus were more active with decreasing Search Risk both in the Search-First condition but also in the Contraband-First condition, although somewhat less significant. In comparison, for the Contraband Risk (associated with the knowledge/recklessness distinction), whereas in the Search-First condition Contraband Risk was robustly positively correlated with a whole range of areas, namely dorsomedial prefrontal cortex (dmPFC), bilateral insula, bilateral TPJ, middle temporal gyrus, DLPFC, and cingulate cortex ( $P < 0.05$ , FWE corrected), for the Contraband-First condition no areas appear, even at a very lenient threshold ( $P < 0.01$ , uncorrected). Thus, the order in which information was presented also had a strong effect on the brain activations associated with Contraband Risk.

Note that, although not identical, there was some overlap between the areas negatively correlated with Search Risk and the areas involved in distinguishing knowing vs. reckless (e.g., bilateral TPJ, DLPFC). These areas appear even though Contraband Risk and Search Risk were modeled as independent events within the same GLM, indicating that they are independently activated by both Contraband Risk and Search Risk. These areas could be generally involved in signaling risk. However, if this were the case, then they should be positively (and not negatively) correlated with Search Risk. Both the TPJ and the DLPFC have been associated with moral decision-making (26, 29). It may be that, as the Search Risk decreases and it becomes more easy to carry contraband across borders, the choice is less of a risky one (how likely am I to get caught?) and more a moral one (should I do it?). Alternatively, it may also well be that these areas are specifically involved in signaling certainty, be it knowing that there was contraband in the suitcase or that they would not be searched. An interesting future study to tackle this issue would be to have participants do a task in which both the risks and potential rewards are similar to the ones adopted in this experiment, but in which there was no legal/contraband-carrying cover story; hence there would be no potential moral dilemma.

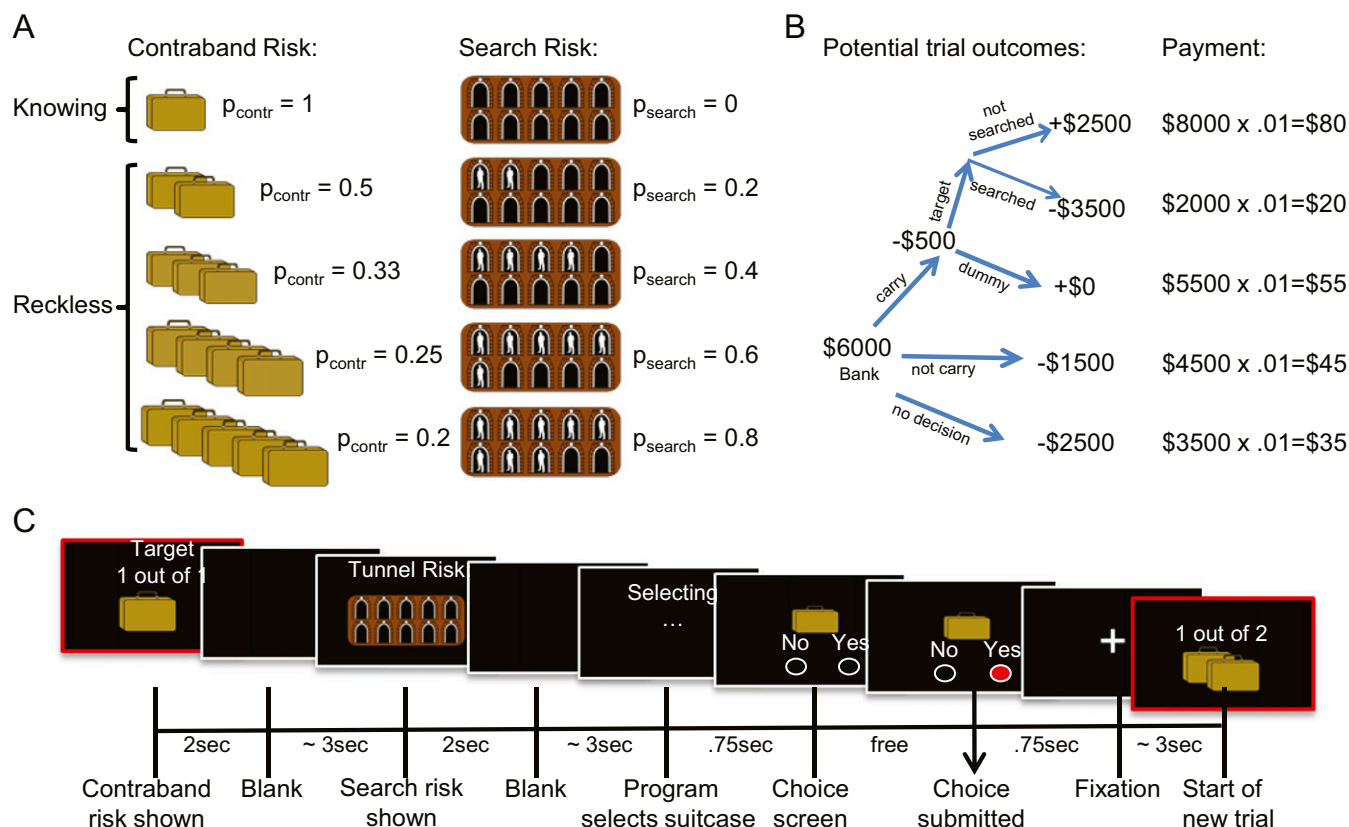
**Prediction Using Different Points in Time Within a Trial.** We chose to model knowledge vs. recklessness at the time the contraband risk is revealed to have a “cleaner” knowledge vs. reckless brain state, and so maximize our chances of observing these two brain states should they exist. To analyze whether the prediction capability of the EN would hold when other times are used, we fitted new models in which knowledge ( $P_{\text{contr}} = 1$ ) and recklessness ( $0 < P_{\text{contr}} < 1$ ) were modeled either at the time Search Risk was being presented or at the time choice was being submitted. For the model comparing knowledge ( $P_{\text{contr}} = 1$ ) vs. recklessness ( $P_{\text{contr}} = 0.33$ ) using the time in which Search Risk was presented, the EN model did not perform better than chance: for the Contraband-First condition, we obtained a mean AUC = 0.475 ( $P_{\text{perm}} = 0.54$ ) and CCR = 37.2% ( $P_{\text{perm}} = 0.67$ ); and the

Search-First condition had a mean AUC = 0.384 ( $P_{\text{perm}} = 0.99$ ) and a CCR = 35.7% ( $P_{\text{perm}} = 0.85$ ). If we use the times in which choice was submitted, the model also does not perform better than chance: for the Contraband-First condition, there was a mean AUC = 0.448 ( $P_{\text{perm}} = 0.62$ ) and CCR = 38.3% ( $P_{\text{perm}} = 0.51$ ); and for the Search-First condition, mean AUC = 0.5 ( $P_{\text{perm}} = 0.42$ ) and CCR = 42.8% ( $P_{\text{perm}} = 0.34$ ). During those times, the brain may be more engaged in processing the current information (search risk or making a decision), and the information about knowledge and recklessness may not be as salient. Thus, the maximum predictability of the model was achieved when modeling the results at the time the contraband risk is presented.

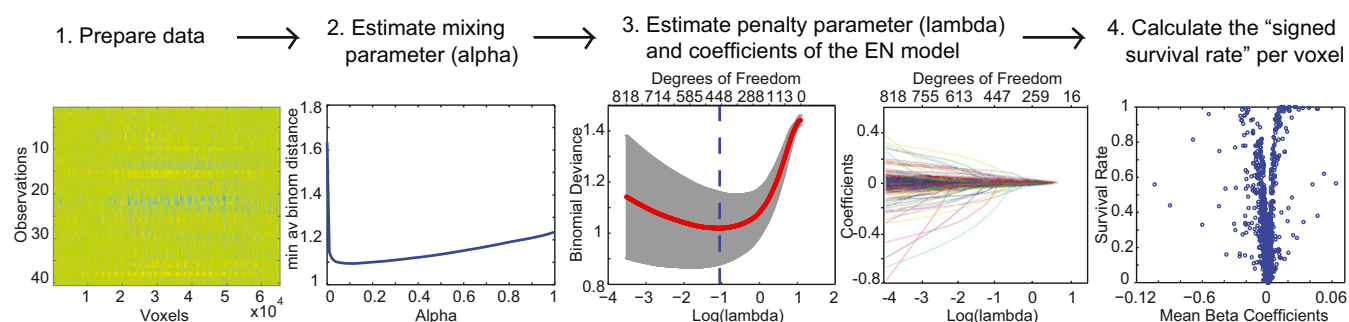
**Predicting Choice.** We can use the EN model approach to try to predict participant's choice (decision to carry or not carry the suitcase) based on their brain data at the time the decision is submitted. We found that, for both conditions, the EN model was able to predict, with high accuracy, the decision of the participant. For the Search-First condition, the EN model had a mean AUC = 1 and a correct classification rate of 99.8% ( $P_{\text{perm}} = 0$ ); and for the Contraband-First condition, the EN classifier had a mean AUC = 0.968 and a mean CCR = 90.9% ( $P_{\text{perm}} = 0$ ). Doing a similar model but in which we try to predict the participant's choice based on brain data at the time the choice screen is first shown (i.e., before they press a button signaling their choice) also yields high performance results: the Search-First condition had a mean AUC = 1 and CCR = 100% ( $P_{\text{perm}} = 0$ ), and the Contraband-First condition had a mean AUC = 1 and CCR = 99.4%. Although these very high performance accuracies are, at least in part, likely due to differential motoric activations (as participants had to press a button in the right hand to say yes and in the left hand to say no), these results serve as a good EN tool validation, showing that it is possible to use the EN regression in both conditions (Search-First and Contraband-First) to distinguish, with very high accuracy, brain states belonging to two different scenarios.

**Brain Areas Specifically Associated with Contraband Risk, Expected Value, and Variance in Reward.** To understand whether the brain pattern we observed used in distinguishing knowing vs. reckless was specifically associated with awareness of Contraband Risk (probability of carrying contraband), or whether it was just related to Expected Value or Variance in Reward, we performed a GLM modeling independently Contraband Risk, Expected Value, and Search Risk. We found that higher Contraband Risk (higher probability) remained positively associated with increased activations in the dmPFC, middle and anterior cingulate cortex, bilateral middle temporal gyrus, bilateral TPJ, and bilateral anterior insula; and negatively associated with bilateral activations in the occipital cortex ( $P < 0.05$ , FWE corrected; Fig. S4A). There were no areas surviving correction for multiple comparisons for higher Variance in Reward, and just the right TPJ was significantly correlated with lower Variance in Reward ( $P < 0.05$ , FWE corrected; Fig. S4B). For Expected Value, the “traditional” areas appeared, namely ventromedial prefrontal cortex (vmPFC) and ventral striatum ( $P < 0.05$ , FWE corrected; Fig. S4C). Both Contraband Risk and Expected Value seem to, independently, activate the superior temporal gyrus, TPJ, and part of the medial PFC (although mainly nonoverlapping areas; Fig. S5). Thus, although some brain areas were activated by several factors, the brain pattern chosen by the EN regression to distinguish knowing vs. reckless seems to be more specifically associated with Contraband Risk.





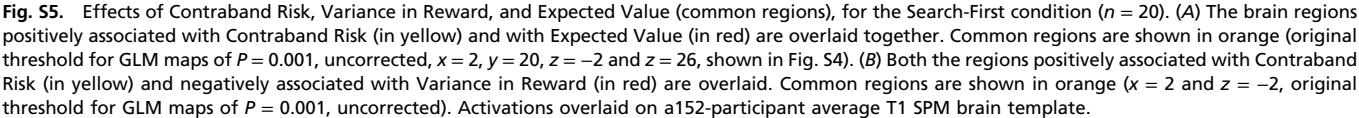
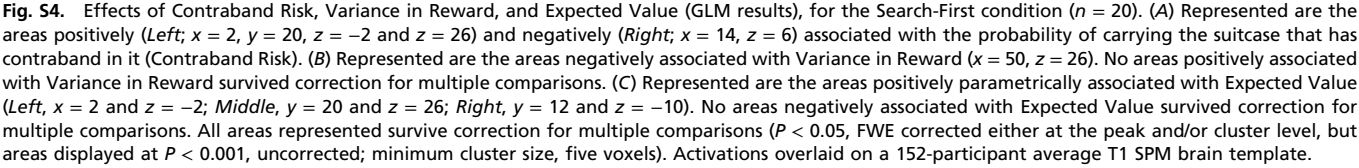
**Fig. S1.** Experimental design. (A) A display of the different scenarios participants were exposed to. (Left) Probability of carrying contraband: Participants were presented with one to five suitcases and asked whether they wanted to carry the suitcase. Only one of the suitcases was the “target” suitcase, supposedly containing contraband, and the other ones were dummies. Hence, when only one suitcase was shown, there was certainty that it was the target suitcase containing contraband (i.e.,  $p_{\text{contr}} = 1$ ). This corresponds to the knowing situation. As the number of suitcases shown increases, there is a lower probability that the person will carry the target suitcase ( $p_{\text{contr}} = 0.5, 0.33, 0.25$ , or  $0.2$  of having contraband in the suitcase). All of these other situations (with probability of carrying the target suitcase lower than 1) correspond to a reckless situation. (Right) Different potential search risk levels. This risk represents the probability of being searched by a “guard.” If the participant is searched and has the target suitcase, he or she incurs a big penalty. The proportion of tunnels with guards indicates the search risk level. (B) Schematic display of the potential decisions and corresponding outcomes that can occur in a given trial. (C) Sequence of events shown to the participants in a typical trial. One-half of the participants ( $n = 20$ ) were shown the contraband risk first, and then the search risk (Contraband-First group), and the other half of the participants were shown the search risk first (Search-First group).

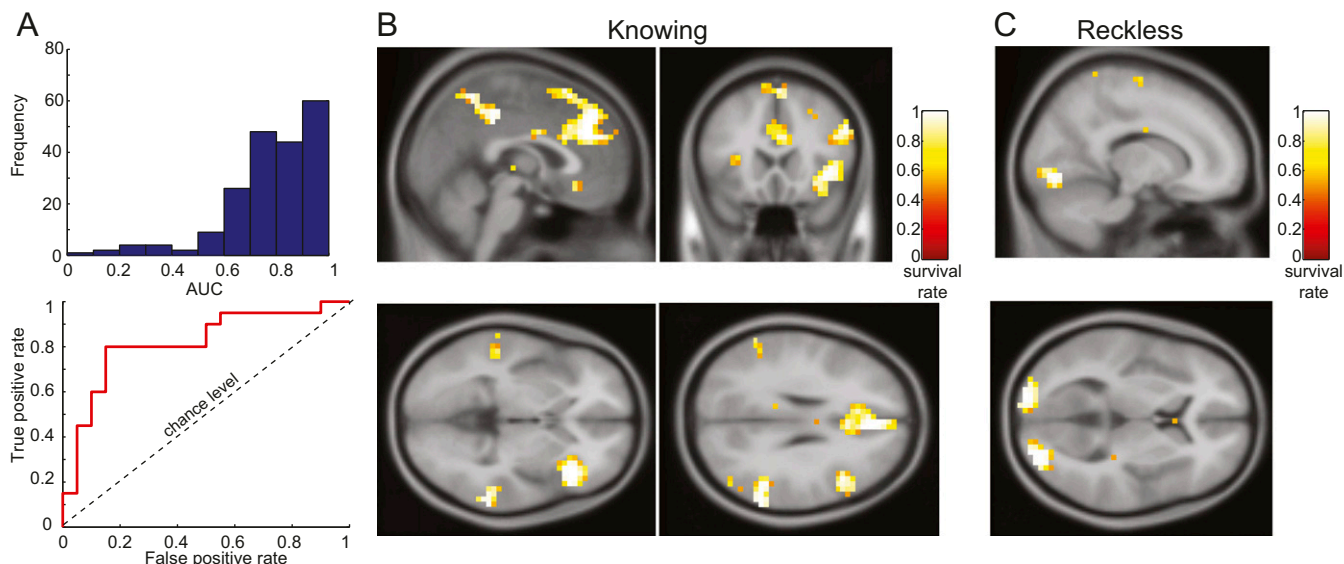


**Fig. S2.** Schematic representation of how the elastic-net (EN) model is implemented. Step 1: Preparing the matrix of data that is going to be used for the EN. For each participant, a simple general linear model (GLM) is first run, which includes as events the occasions in which they see “one suitcase” and also when they see “five suitcases.” The corresponding beta maps are extracted and converted in rows, so that each participant is represented in two rows, one for “one suitcase” (knowing) and one for “five suitcases” (reckless). The length of each row corresponds to the total number of voxels present (e.g.,  $\sim 65,000$ ), minus the columns associated with NaN values (corresponding to nonbrain data). Step 2: Using the matrix obtained in 1, a grid search over  $\alpha$  is performed, and the  $\alpha$  that minimizes the average minimum binomial distance is chosen. Step 3: The EN is fitted over a range of  $\lambda$  values, and for each  $\lambda$  a series of coefficients is obtained (one coefficient per voxel). Then, the  $\lambda$  that minimizes the binomial deviance is chosen (minimum  $\lambda$ ). For this  $\lambda$  ( $\lambda_{\text{min}}$ ), we extract the corresponding coefficients and register which voxels “survived” (i.e., had coefficients different from zero), using fivefold cross-validation. This procedure is then repeated 40 times, for a total of 200 runs (fivefold cross-validation  $\times$  40 repetitions). Step 4: After all of the iterations are done, the survival rate is calculated for each voxel, multiplied by the sign of the average coefficient value over 200 CV runs, leading to a “signed survival rate” value per voxel. See *SI Materials and Methods* for details.

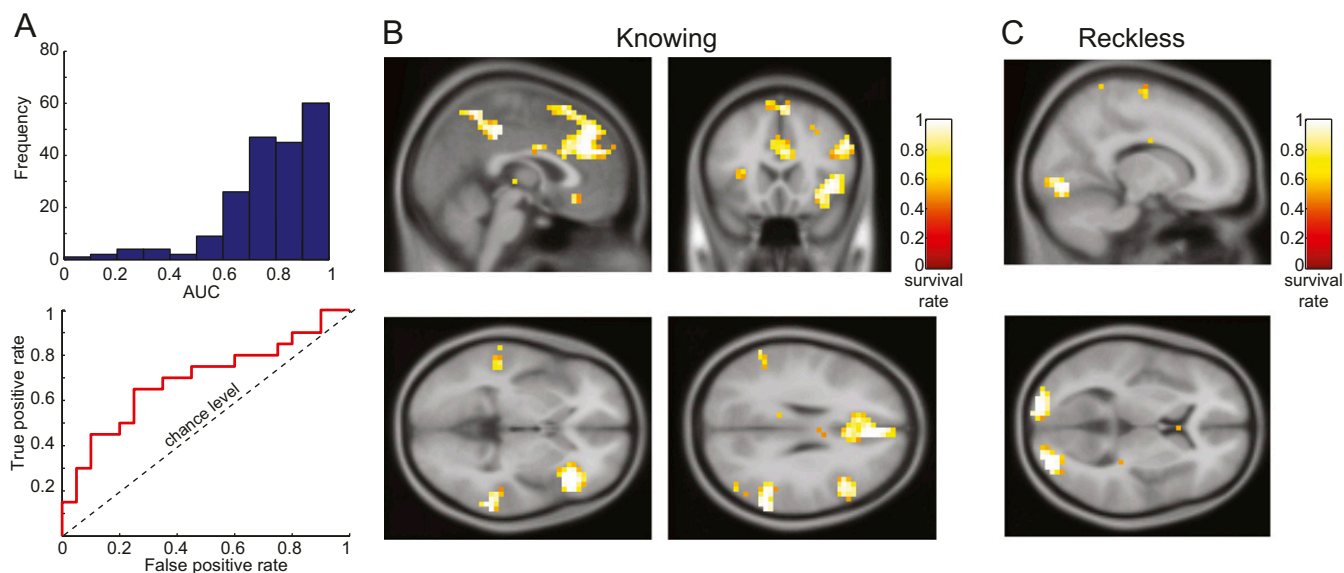








**Fig. S6.** The knowing/reckless distinction, extracting out parametric effects associated with the probability of being caught. EN results when using as a baseline the betas from a GLM in which the effects of probability of being caught (searched and with target suitcase) were modeled separately and not included in the EN model (Search-First condition,  $n = 20$ ). The model led to higher-than-chance prediction accuracies, with a CCR = 71.8% and an AUC = 0.792 ( $p_{\text{perm}} = 0.005$ ). (A, *Top*) Distribution of cross-validated areas under the curve (AUCs) averaged over 200 values. (*Bottom*) Example of one ROC curve. (B) Voxels more associated with a knowing situation are shown (negative surviving-rate voxels;  $x = 2$ ,  $y = 20$ ,  $z = -2$  and  $z = 26$ ). (C) Voxels more associated with a reckless situation are presented (positive surviving-rate voxels,  $x = 14$ ;  $z = 6$ ). Each voxel's (signed) survival rate is overlaid on an axial section of a 152-participant average T1 SPM brain template higher (with a minimum survival rate for the cluster's peak voxel of 0.5).



**Fig. S7.** The knowing/reckless distinction, extracting out parametric effects associated with the probability of obtaining the highest reward. EN results when using as a baseline the betas from a GLM in which the effects of probability of obtaining the highest reward (not searched and with target suitcase) were modeled separately and not included in the EN model (Search-First condition,  $n = 20$ ). The model still retained its high prediction ability (CCR = 71.7% and AUC = 0.792,  $P_{\text{perm}} = 0.005$ ). (A, Top) Distribution of cross-validated areas under the curve (AUCs) averaged over 200 values. (Bottom) Example of one ROC curve. (B) Voxels more associated with a knowing situation are shown (negative surviving-rate voxels;  $x = 2$ ,  $y = 20$ ,  $z = -2$  and  $z = 26$ ). (C) Voxels more associated with a reckless situation are presented (positive surviving-rate voxels,  $x = 14$ ;  $z = 6$ ). Each voxel's (signed) survival rate is overlaid on an axial section of a 152-participant average T1 SPM brain template higher (with a minimum survival rate for the cluster's peak voxel of 0.5).

Coefficient	Estimate	SE	z score	Pr(> z )
Intercept	-4.1977	0.2333	-17.992	<2e-16
Search Risk	5.3727	0.3985	13.482	<2e-16
Contraband Risk	-0.4321	0.4967	-0.870	0.38436
Search*Contraband	2.8369	0.8799	3.224	0.00126
Search*Condition	-0.1337	0.2640	-0.507	0.61238
Contraband*Condition	-0.1235	0.4776	-0.258	0.79603
Search*Contraband*Condition	2.5885	1.0680	2.424	0.01536

Logistic regression of response (positive coefficient means for increased probability of rejecting carrying the suitcase) on Search Risk (0, 0.2, 0.4, 0.6, 0.8), Contraband Risk (0, 0.2, 0.25, 0.33, 0.5, 1), the interaction of Search and Contraband Risk (Search\*Contraband), interaction of Search Risk and Condition (Condition 1 is contraband risk seen first, and Condition 2 is search risk seen first), interaction of Contraband Risk and Condition, and the three-way interaction of Contraband Risk, Search Risk, and Condition.

Table S2. Surviving voxels

Conditions	H	Area	Coordinates			Survival rate
			x	y	z	
Search-First knowing	L/R	Dorsomedial prefrontal cortex	2	40	46	1
	L/R	Anterior cingulate cortex	6	36	26	1
	R	Insula	42	24	2	1
	R	Temporoparietal junction	58	-48	30	1
	L	Temporoparietal junction	-58	-48	30	1
	R	Middle temporal gyrus	58	-40	2	1
	L/R	Middle cingulate cortex	6	-36	38	1
	R	Dorsolateral prefrontal cortex	46	16	26	0.95
	L/R	Medial orbitofrontal cortex	2	28	-14	0.85
	L	Middle temporal gyrus	-54	-32	-2	0.81
reckless	L	Insula	-30	20	10	0.67
	R	Visual cortex	26	-88	6	1
	L	Visual cortex	-18	-96	2	0.99
	R	Supplementary motor area	14	-12	70	0.73
Contraband-First knowing						
	R	Temporoparietal junction	50	-60	46	0.26
reckless	L	Visual cortex	-30	-88	2	0.46
	R	Orbitofrontal cortex	38	32	-14	0.30
	R	Visual cortex	34	-84	14	0.28

Represented are the *x*, *y*, *z* coordinates for (one of the) peak voxels within each brain area used by the EN regression to predict knowing vs. reckless. H, Hemisphere. Coordinates are listed in standard Montreal Neurological Institute space.