Supporting Information

Bhatt et al. 10.1073/pnas.1009625107

SI Methods

Subject Information. Subjects were recruited from a large database maintained by the Human Neuroimaging Laboratory (HNL) at Baylor College of Medicine (BCM). Most subjects were affiliated with the Texas Medical Center or Rice University in various ways. Subjects were recruited in pairs and kept separate throughout the task, so the experiment was completely anonymous. Subjects were informed that they were playing a real person and seemed to believe this to be the case.

After filling out a standard screening form and giving consent, subjects were given written instructions on the task. After reading the instructions they were walked through the task again verbally and asked to answer a few questions about hypothetical bargaining rounds to make sure that they fully understood the task. If any of these questions was answered incorrectly, the mistake was explained and another scenario presented.

After both subjects were fully instructed, they were loaded into the scanner and participated in 60 rounds of the bargaining task. At the end, they were told their total earnings and given an open-ended debriefing sheet on which they were asked to describe, in their own words, how they had approached the task.

Demographics, including age, sex, and socioeconomic status for all 76 subjects, are provided in Table S1 according to group. In addition to those three factors we assessed 29 subjects (10 incrementalists, 9 conservatives, and 10 strategists) on a standard IQ test [Cattel's Culture Fair IQ test, form 3a (1)]. The only significant demographic effect among groups according to one-way ANOVA was for age: strategists are significantly younger than other subjects. Conservative and strategist IQs were both significantly higher than average, whereas incrementalist IQs were not.

Behavioral Analysis. For each of the 76 buyers in the sample we classified behavior according to the last 30 periods of the game. We focused on the latter half of the experiment because several subjects changed strategies significantly between the first and second halves of the experiment. Of particular interest, four subjects in this analysis became strategists behaviorally midway through the experiment. Not only did their behavior shift significantly, but also upon debriefing these subjects explicitly mentioned changing from Nash behavior (no correlation between suggestion and value) to strategist behavior (a negative correlation between suggestion and value). Although this shift in itself is interesting, it occurred in too few subjects to do systematic analysis (a similar evolution of strategic sophistication is shown parametrically in ref. 2). We therefore concentrated on the latter half of the experiment when all subject behavior had stabilized. Results from the behavioral regressions in both early and late rounds are shown in Table S2. Note that between early and late rounds, the slope coefficients ["information revelation" (IR) values] for both incrementalist and strategists generally increased in magnitude, and R^2 values increase.

Cluster analysis. We performed a regression of buyers' suggestions on their private values from the second half of the experiment. This yielded three descriptive strategy parameters for each buyer: the slope, intercept, and fit (R^2) of the regression. We normalized these three statistics across subjects by subtracting means and dividing by SDs. Clusters were then identified using the *k*-means algorithm (3). We found that cluster assignment only changed for one subject when intercepts were included (that assignment went from incrementalist to conservative). Therefore, the results in the text are clustered using the *k*-means method on only slope and fit.

Cognitive hierarchy model. The cluster analysis was a bottom-up, data-driven analysis that did not address how subjects might arrive at behavior. To complement this analysis, we created a model based on Camerer et al.'s cognitive hierarchy (CH) model (4) that addresses player incentives, the structure of the game, and most importantly posits an algorithm for how individuals might form beliefs in this game.

In the cognitive hierarchy model (Fig. S1), players perform different numbers of "steps of thinking" to form beliefs. Zero step (or "level-0") thinkers behave naïvely—randomly in the original formulation—and lack a model of how other players will behave. Level-1 thinkers assume that they are playing level-0 players and best respond to naïve behavior. Level-n thinkers think that they are playing a mixture of players from all of the n-1 levels below them and best respond to that mixture.

In our model we assume that level-0 buyers have a fixed type, α, distributed u(0,1). Level-0 buyers will thus send suggestions according to:

$$s = \min(10,\ \max(1,\ [\alpha v + \varepsilon])),\ \varepsilon \sim \mathrm{N}(0, \sigma^2).$$

Level-0 sellers are assumed to be naïve and respond to a buyer suggestion s according to:

$$p = \min(10,\ \max(1,\ [s + \varepsilon])), \varepsilon \sim \mathrm{N}(0, \sigma^2).$$

Here $[x]$ is the nearest integer function. We add the max/min operations to account for the fact that both price and suggestion must be integer valued and between 1 and 10. Let F denote the cumulative normal distribution function with mean 0 and variance σ^2. Then these assumptions translate to the conditional distributions:

$$P_{0-buyer}(s \mid v, \alpha) = F(s + .5 - \alpha v) - F(s - .5 - \alpha v) \quad \text{for } 1 < s < 10,$$

$$P_{0-buyer}(1 \mid v, \alpha) = F(1.5 - \alpha v) \text{ and } P_{0-buyer}(10 \mid v, \alpha) = 1 - F(9.5 - \alpha v).$$

$$P_{0-seller}(p \mid s) = F(p + .5 - s) - F(p - .5 - s) \text{ for } 1 < p < 10,$$

$$P_{0-seller}(1 \mid s) = F(1.5 - s) \text{ and } P_{0-seller}(10 \mid s) = 1 - F(9.5 - s).$$

Intuitively, buyers suggest a price that is a fraction α (<1) of their value, plus noise. This represents "shaving" bids, a behavior commonly observed in auctions of this type (5).

We assume that level-1 buyers respond optimally to level-0 sellers, with some noise according to the softmax distribution:

$$P_{1-buyer}(s \mid v) = \frac{\exp(\lambda \pi_{1-buyer}(s, v))}{\sum_{s'} \exp(\lambda \pi_{1-buyer}(s', v))}.$$

In this expression $\pi_{1-buyer}(v, s)$ is the expected payoff given value v and suggestion s if you are a level-1 buyer:

$$\pi_{1-buyer}(v, s) = \sum_{p<v} P_{0-seller}(p \mid s)(v - p).$$

The level-1 buyer's probability of sending s is decreasing in v, with the distribution shifting downward as v increases.

Level-1 sellers are assumed to also respond optimally to level-0 buyers. However, they have the extra computational challenge of updating their prior distributions over the value of α, so their choice of price depends on the current suggestion and also on the entire history of suggestions (and on the sensible inference of how much information about value the suggestions have generally implied). They respond according to the distribution:

$$P_{1-seller}(p\,|\,\{s_t\}) = \frac{\exp(\lambda\pi_{1-seller}(p\,|\,\{s_t\}))}{\sum_{p'}\exp(\lambda\pi_{1-seller}(p'\,|\,\{s_t\}))}$$

$$\pi_{1-seller}(p\,|\,\{s_t\}) = p\cdot P(v\geq p) = p\sum_{v\geq p}P_{0-buyer}(v\,|\,\{s_t\}).$$

We find the distribution $P_{0-buyer}(v,\{s_t\})$ in two steps. First we determine the posterior distribution over α at time t, using Bayesian updating iteratively:

$$P(\alpha\,|\,\{s_t\}) = \frac{P(s_t\,|\,\alpha)P(\alpha\,|\,\{s_{t-1}\})}{\int_0^1 P(s_t\,|\,\alpha')P(\alpha'\,|\,\{s_{t-1}\})d\alpha'},$$

where

$$P(s\,|\,\alpha) = \sum_{v=1}^{10}P_{0-buyer}(s|v,\alpha).$$

We use this posterior over α to find a posterior over values:

$$P_{0-buyer}(v\,|\,\{s_t\}) = \frac{\int_0^1 P(v\,|\,\alpha,s_t)P(\alpha\,|\,\{s_t\})d\alpha}{\sum_{v=1}^{10}\int_0^1 P(v'\,|\,\alpha,s_t)P(\alpha\,|\,\{s_t\})d\alpha},$$

where

$$P(v\,|\,\alpha,s) = \frac{P_{0-buyer}(s\,|\,v,\alpha)}{\sum_{v'=1}^{10}P_{0-buyer}(s\,|\,v',\alpha)}.$$

We simplified these computations by approximating the initial uniform distribution over α with a discrete distribution whereby α took the values {0, 0.01, 0.02, . . ., 0.99. 1} with equal probability.

Finally, level-2 buyers respond optimally to a 50/50 mixture of level-0 and level-1 sellers. Importantly, level-2 buyers anticipate how their suggestions will change a level-1 seller's posterior over α and how these changes will affect payoffs in the next period. (Because we are assuming limited, hierarchical reasoning throughout, we only have level-2 sellers project one period into the future rather than considering the entire experimental run. We consider only one period of forecasting because adding more does not significantly change predicted choices but does become too computationally taxing to estimate. Time to estimate the level-2 model grows exponentially with the number or periods ahead considered.) To get probabilities of buyer's suggestions we plug expected payoffs given these beliefs into a softmax function as we did for level-1 buyers and sellers:

$$P_{2-buyer}(s_t\,|\,v,\{s_{t-1}\}) = \frac{\exp(\lambda\pi_{2-buyer}(s_t\,|\,v,\{s_{t-1}\}))}{\sum_{s'}\exp(\lambda\pi_{2,buyer}(s'\,|\,v,\{s_{t-1}\}))}.$$

$$\text{Here } \pi_{2-buyer}(s_t\,|\,v,\{s_{t-1}\}) = \text{Current Payoff} + \text{Future Payoff where :}$$

$$\text{Current Payoff} = .5\Big(\sum_{p<v}P_{1-seller}(p\,|\,\{s_t\})(v-p)\Big) + .5\Big(\sum_{p<v}P_{0-seller}(p\,|\,s_t)(v-p)\Big), \text{and}$$

$$\text{Future Payoff} = \frac{\sum_{v_{t+1}}\max_{s_{t+1}}\Big(\sum_{p<v_{t+1}}(.5P_{1-seller}(p\,|\,\{s_{t+1}\}) + .5P_{0-seller}(p\,|\,s_{t+1}))(v-p)\Big)}{10}$$

Predicted behavior for level-1 sellers in period 2 under different suggestion histories (s = 1 in period 1 vs. s = 2 in period 2) is shown in Fig. S2 along with predicted level-2 buyer behavior in period 1. Note that at low values, level-2 buyers will tend to send high suggestions because this will improve their period 2 expected payoffs, whereas at high values they send low suggestions to collect current rewards. Note that this behavior corresponds to the strategist cluster in the regression-based analysis.

We estimated the maximum likelihoods for level-0, level-1, and level-2 behavior for each of the 76 buyers in the sample, concentrating once again on the second half of the experiment as we did in the cluster analysis. We found that there were problems separately attributing noise in the models to the λ and σ parameters (i.e., identifying parameters), so we fixed σ = 2. Because each of the models then has only one free parameter—α for level-0, λ for levels 1 and 2—we compared the three likelihoods directly to classify each buyer as level 1, 2, or 3. Parameters were optimized over the following ranges: α in [0, 1], λ in [0.5, 8]. Table S3 summarizes these results. Log likelihoods were not reported if they did worse than random (log likelihood of −69). Two subjects, AQ-1 and DY-1, were not classified as level 0, 1, or 2 because the random model does better than all three. Fourteen of the 16 strategists in our sample were classified as level 2. The remaining strategists were classified as level 1. All incrementalists were classified as level 0; however, the CH model tended to classify conservatives as level 0 players with low αs rather than as level 1.

Functional MRI Analysis. ***Other activations of interest.*** In addition to the results shown in the main text, there were several other activations of interest that were suggestive but did not pass any correction for multiple comparisons. There was a significant activation in the right retrosplenial cortex when we regressed activity at decision time on IR. Activation and time courses separated by behavioral type are shown in Fig. S3*A*. Although the activation was detected at the time of decision, time course analysis reveals that the differences in activation begin at trial onset and continue through the postdecision epoch. Although the functions of this area are less well understood than those of the areas highlighted in the main text, we believe that this activation is indicative of the prospective thinking involved in strategist thinking (6–8).

Incrementalists showed decreased activation in the right caudate compared with the other two groups (boxcar over trial; Fig. S3*B*). This area has been linked to perceptions of reward, particularly when those rewards are seen as the result of specific actions (9). Decreased activity in this area in the incrementalists may reflect their relatively passive, honest strategy. In addition, analysis of the region of interest (ROI) reveals that only strategists show consistent modulation of caudate activity by value.

We performed an additional set of analyses using R^2 as the between-subject regressor and found two activations of interest. First, we found that bilateral posterior insula correlated positively with R^2 at trial onset (Fig. S4*A*). Second, we found a significant negative interaction between suggested price and R^2 in the left posterior superior temporal sulcus (STS). ROI analysis of the latter activation indicated that this activation was driven by a negative correlation between activity and suggested price in the incrementalists that was largely absent in the other two groups. Examination of the time series showed that incrementalists seemed to have a secondary activation in the area at decision after sending low, but not high, suggestions (Fig. S4*B*).

In all of the analyses, the strategists drove most of the activations. We found very few areas where activity was actually higher for the incrementalists or conservatives. Two such activations in the dorsal anterior cingulate cortex (ACC) and middle paracingulate at decision ended up actually being the timing artifacts of larger activations extending over the trial, and activation peaks did not prove to be significantly different (Fig. S5). The interpretation of this activation as a timing artifact is supported by the facts that incrementalists have somewhat longer response times and that although the activation in question begins at trial onset and persists through the trial, no between-subject differences are found using the boxcar regressions that account for these differences in response times. These areas are also apparent in the analysis using R^2 as a between-subject regressor at decision (Table S4*b*; clusters of interest are italicized). The absence of between-group differences in these areas is notable given Coricelli and Nagel's findings that medial prefrontal cortex activity correlates with depth of reasoning in the p-beauty contest (10).

Activation Tables. Activations from the relevant between-subjects contrasts on the main effects are shown in Table S4 *a–c*. Relevant interactions are shown in Table S4*d*. Cluster corrections are shown at $P < 0.001$ (uncorrected) and cluster size (k) at least 5. Between-group t tests and regression on IR often showed very similar areas of activation, and all three of the areas discussed in the main text were significant at the 0.001 level in both analyses. Areas discussed in the main text are highlighted in bold; areas highlighted in the supporting information are italicized.

1. Cattell RB (1963) Theory of fluid and crystallized intelligence: A critical experiment. *J Educ Psychol* 54:1–22.
2. Chong J, Camerer C, Ho T (2006) A learning-based model of repeated games with incomplete information. *Games Econ Behav* 55:340–371.
3. MacQueen J (1967) Some methods for classification and analysis of multivariate observations. *Proc Fifth Berkeley Symp Math Stat Prob* 1:281–297.
4. Camerer CF, Ho T, Chong J (2004) A cognitive hierarchy model of games. *Q J Econ* 119: 861–898.
5. Camerer CF (2003) *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton Univ Press, Princeton, NJ), pp 187–195.
6. Addis DR, Wong AT, Schacter DL (2007) Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. *Neuropsychologia* 45:1363–1377.
7. Luhmann CC, Chun MM, Yi DJ, Lee D, Wang XJ (2008) Neural dissociation of delay and uncertainty in intertemporal choice. *J Neurosci* 28:14459–14466.
8. Vann SD, Aggleton JP, Maguire EA (2009) What does the retrosplenial cortex do? *NatRev Neurosci* 10:792–802.
9. Tricomi EM, Delgado MR, Fiez JA (2004) Modulation of caudate activity by action contingency. *Neuron* 41:281–292.
10. Coricelli G, Nagel R (2009) Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proc Natl Acad Sci USA* 106:9163–9168.

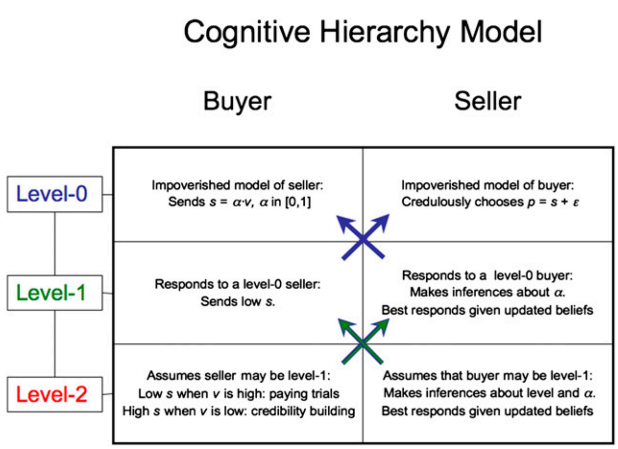

Fig. S1. Cognitive hierarchy model. In the model, level-0 players are assumed to be naïve. In this case, naïve buyers use a simple linear suggestion strategy, and level-0 sellers are credulous. Level-n players believe that they are playing a lower-level player. Thus, a level-1 player believes that she is playing a naïve level-0 player with certainty. On the other hand, a level-2 player has a prior distribution over opponent type (level-0 or level-1) and best responds given this prior distribution.

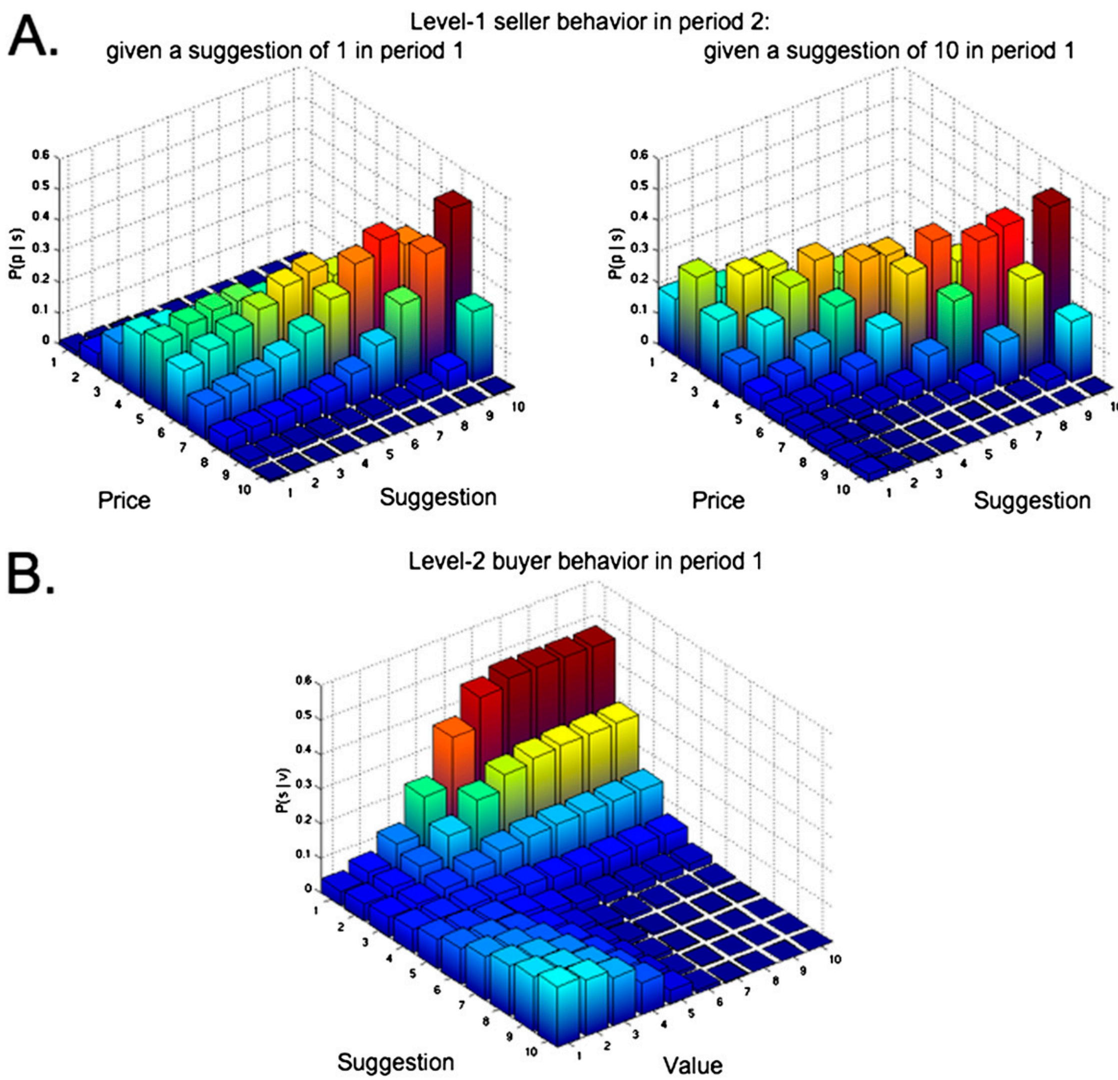

Fig. S2. Model predictions. (*A*) Predicted level-1 seller behavior in period 2 given different suggestions in period 1. On the left, note that when the period 1 suggestion is 1, level-1 sellers will be relatively insensitive to low suggestions, choosing an average price of 5 given a suggestion of 1. However, when the suggestion in period 1 is 10, sellers are far more likely to respond to a low suggestion with a low price (behavior shown on the right). (*B*) Predicted level-2 buyer behavior in period 1. Given the behavior shown in *A*, level-2 buyers are likely to choose a high (>5) suggestion in period 1 if their value is low, to increase possible future payments.

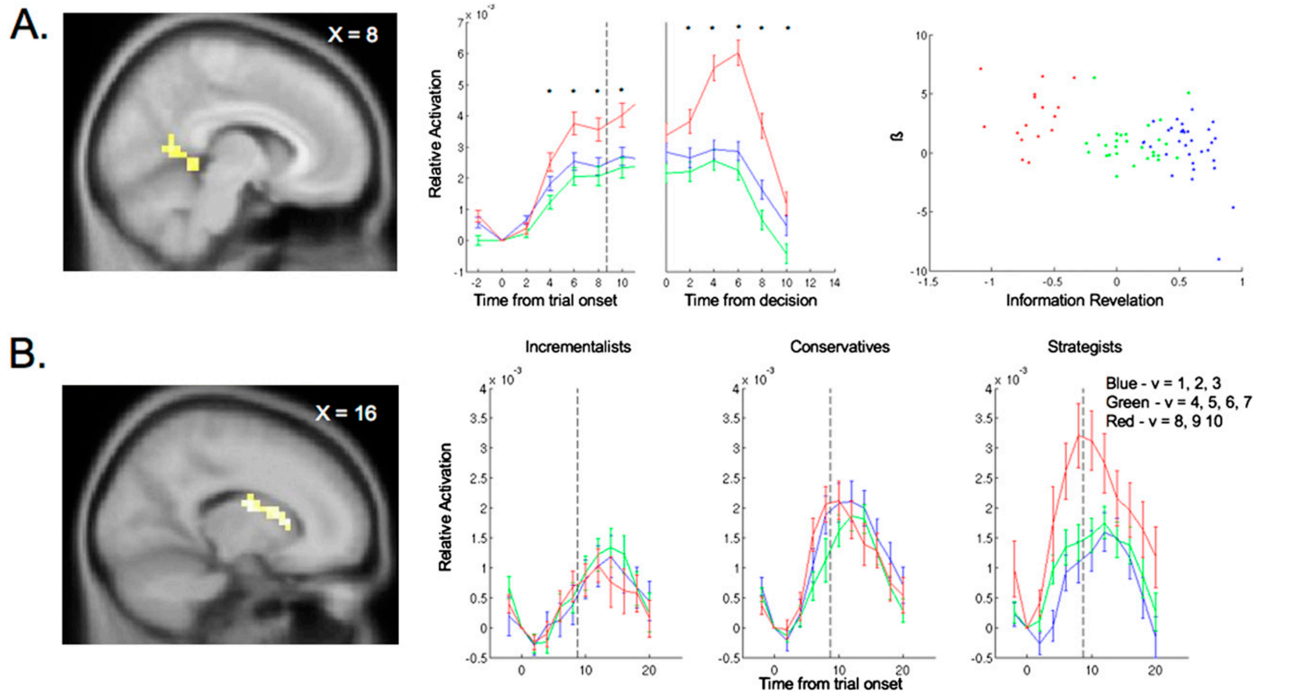

Fig. S3. Other activations of interest. (*A*) Strategists differentially activate the retrosplenial cortex at the moment of deception. Regression of neural activity at decision on information revelation. *Left:* Retrosplenial activation at decision time. Peak voxels at (8, −56, 8) and (8, −40, 0), $k = 21$, cluster-level $P = 0.088$ (corrected). *Center:* Time courses in retrosplenial cortex by group. *Right:* Scatter plot of general linear model coefficients vs. IR. (*B*) Activity along the right caudate over entire trial. Right caudate was found to be less active on average in incrementalists than in conservatives or strategists when using the boxcar regressor. Cluster shown here (*Left*) is thresholded at $P < 0.002$; activation remains at $P < 0.001$. Examination of the time series broken down by the buyer's value (*Right*) shows that not only is overall activity increased in the conservative and strategist groups, but only the strategists show consistent modulation of activity by value.

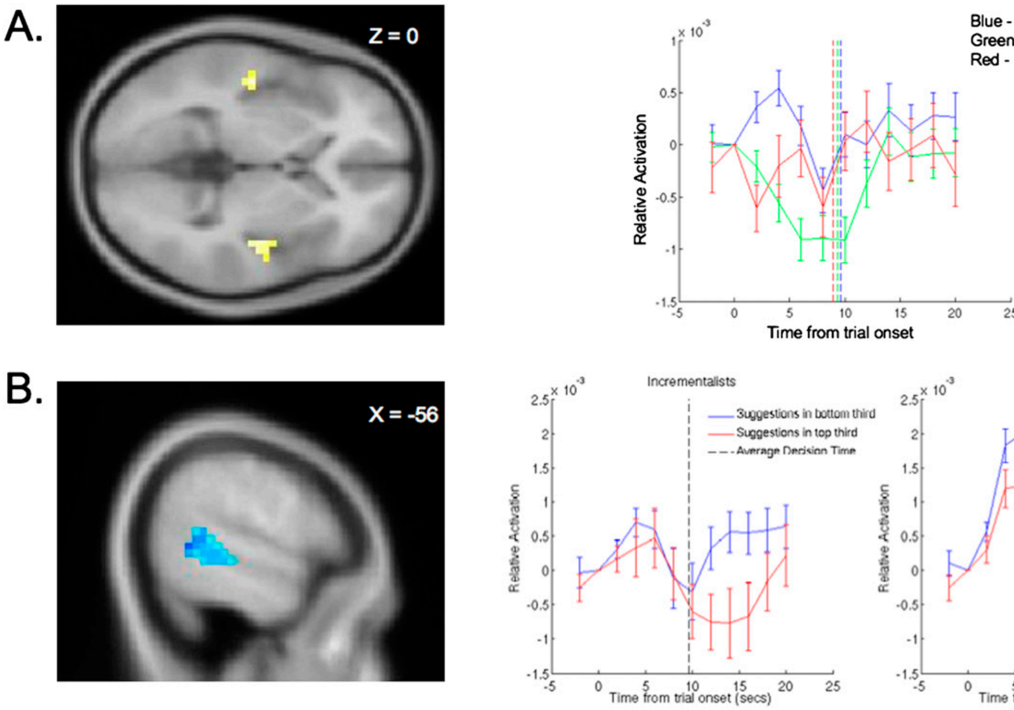

Fig. S4. Activations correlated to R^2. (*A*) Bilateral posterior insula activity correlates with R^2 at trial onset. Examination of the time series (*Right*) indicates that this may be driven by a deactivation in the area in conservatives. Average decision times for each of the three groups are shown by dashed lines. (*B*) Activity in left posterior STS at decision time correlates negatively with the interaction between R^2 and suggested price. The time series analysis (*Right*) indicates that this is caused by a secondary activation in incrementalists in the area after they send a low suggestion, which is absent when they send high suggestions. Strategists show no difference in STS activity based on choice, and conservatives show insufficient variance in suggestions to find any relationship. Activations are both shown at $P < 0.001$, uncorrected. Full statistics are in Table S4.

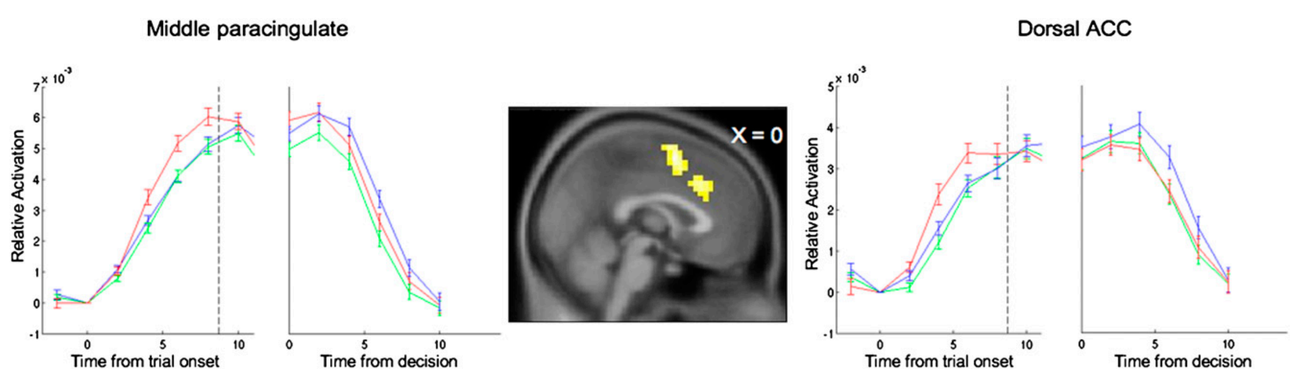

Fig. S5. There was an apparent incrementalist activation of the dorsal ACC and middle paracingulate at subject decision time; however, examination of the time courses shows that this may be driven by slight differences in the timing of an activation at trial onset between groups. Specifically, activity in the area seems to peak slightly earlier for strategists, manifesting as a difference in activation at decision. There were no significant differences among the peaks of these activations.

Other Supporting Information Files

Table S1 (DOC)
Table S2 (DOC)
Table S3 (DOC)
Table S4 (DOC)