NEURAL ACTIVITY DURING BARGAINING WITH PRIVATE INFORMATION

MEGHANA BHATT, TERRY LOHRENZ, READ MONTAGUE, AND COLIN F. CAMERER

Abstract. We monitored subjects' braing activity during a simple bargaining task with cheap talk using fMRI. We found that areas shown to be correlated with cognitive control we active when subjects ignored the cheap talk signals they received, indicating that treating information as credible may be an automatic or "type one", response requiring cognitive effort to override. These automated responses may explain the fact that over-communication leading to greater efficiency is the norm in these sorts of games.

1. Introduction

Strategic interaction lies at the heart of social science. Game theory provides a precise language to study strategic interaction and provides analytical tools for predicting what will happen when people interact. In a game, players choose strategies given information about the strategies and private information about other players. Strategy choices create outcomes which players are assumed to value, in a way that can be expressed by numerical *utility*.

The central complication in analysis of a game is what players believe other players will do. Beliefs, and the sensible strategy choices that result given different beliefs will lead to different outcomes. They key questions in predicting outcomes are first, what characterizes a "sensible" strategy choice, and second, what determines our beliefs.

Economists have traditionally gotten around these questions using equilibrium concepts. A sensible strategy is simply a best response, i.e. expected utility maximizing, response given beliefs. The question of belief is replaced by that of mutual best response. However equilibrium concepts like Nash equilibrium give mixed results as predictors of human behavior (cites, cites, cites).

This is an extremely preliminary working report for presentation at the ASSA meetings, January 2007. Do not quote or circulate without permission but comments are welcome. Bhatt, Camerer: Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, camerer@hss.caltech.edu. Lorenz and Montague: Baylor School of Medicine. Direct correspondence to the first author. This research was supported by a Moore Foundation grant to CFC and ?? to RM. Thanks to Jacob Goeree, Brooks King-Cassas, Antonio Rangel and Joseph Wang for helpful comments.

An important canonical game in economics, political science, biology and other fields is bargaining. The exchange of goods, either for other goods, or for currency is a basic activity that occurs in any human culture. However this sort of game requires both parties to have a good understanding of the desires or the other.

In this paper we study a simple version of this game with "cheap-talk", a limited round of communication with no direct effect on the outcome of the game. A buyer and seller bargain briefly over the price at which the seller will trade an object with zero marginal cost in one period of bargaining. Since the seller's cost is zero, she could sell at any price. The complication is that the buyer's value for the object, v, is drawn from a uniform distribution of integers from 1 to 10. The buyer knows her value but the seller only knows that the value is uniformly distributed.

The first step in the game is that the buyer "suggests" a price s. This suggestion is costless and has no direct payoff consequences. The buyer hears the suggestion s and names a take-it-or-leave-it price p. If the price p is above or equal to the value v ($p \geq v$) the trade takes place at the price p. The seller then earns $p - 0 = p$ and the buyer earns the surplus $v - p$. Notice that the suggestion s does not enter directly into what the buyer and seller earns; it just serves as a communication device which might influence the price the seller demands.

Despite its simplicity, this game is interesting from both an economic view, and from the viewpoint of social neuroscience. In the standard analysis (which assumes self-interest and no computational costs), the seller realizes the buyer's suggestion is always designed to get her (the seller) to name a lower price. The seller should therefore ignore s, and name the price which maximizes her expected profit, either 5 or 6 (both prices are equally profitable). Anticipating the fact that the seller ignores s, the buyer should "babble", in game theory jargon, and choose a value of s which is unrelated to v. However, since the buyer's value is always above the seller's cost (zero), a failure to trade is socially inefficient. This game therefore features a conflict between the social desire to always trade, and the individual desires to get the best price.

In experiments of this general type, buyers generally "overcommunicate" (there is a measurable correlation between the true state and the signal sent) and sellers are sensitive to the buyers' suggestions, contrary to the prediction of theory. The gains from exchange that result from this pattern are higher for the two players together than if there was babbling by the buyer and ignorance of the suggested price by the seller **????**.

Our bargaining game is of interest in social neuroscience too. Compared to other species, humans are highly social and have created complex social architectures for trade, transmission of cultural practices and information, and institutional rules such as laws and social norms. These practices are presumably supported by neural architecture for decoding and creating facial expressions, processing abstract signals of who can be trusted, and creating reasonable "models" of others behavior: both understanding their desires and how those desires are linked to their actions. While there are some fMRI studies of social exchange in games involving trust **??**, and randomization **?**, there are no studies of the sort of strategic communication that results from bargaining.

An important open question is why and when overcommunication occurs: Are one or both sides making a strategic mistake? Are they being prosocial and cooperative, knowing that choosing an informative value of s and responding to it will increase how much they make together? Or is there something even more basic going on, involving how our brains are wired to process communication from another person.

2. Background

We chose to study bargaining because it provides a natural strategic environment for subjects, bargaining over an object is a familiar situation to many adults; and because the incomplete information in the problem, namely the private values for the buyer and sellers, provide a rich environment to explore belief formation. In addition bargaining is a well-understood problem theoretically. Myerson and Satterthwaite discussed the *inherrent* inefficiency of bilateral bargaining using any incentive compatible mechanism**?**.

We chose to study cheap talk in these games for several reasons. First, experiments show that cheap talk will increase the average efficiency of bilateral bargaining games **?**. Second, the core purpose of cheap talk is to manipulate another player's beliefs. Crawford and Sobel **?** characterized the Bayes-Nash equilibria of a general class of sender-receiver games. In these games there is some randomly chosen state of the world that is revealed to the "sender", the sender is then allowed to send some signal s to a "receiver", after seeing this signal the receiver takes an action $a \in A$ that determines the outcome of the game.

They predict that the level of information transmission is proportional to how well alligned the player's preferences are in games similar to the task we study here **?**, but as mentioned above, in bargaining games the buyer's interests are almost directly opposed to the seller. In addition they show that any equilibrium is equivalent to one of the class described below:

- Let R be the possible states of the world and S be the set of possible signals. There exists some function signaling function $\sigma\ R \rightarrow S$ such that the sets $\sigma^{-1}(s)$ create a partition over R.
- The receiver has some response function $\rho\ S \rightarrow A$ such that $\rho(s)$ maximizes the receiver's expected payoff conditional on the fact that the state of the world is an element of $\sigma^{-1}(s)$.

The allows us to use the number of sets in the partition induced by σ^{-1} as a measure of how informative the equilibrium is. For example if there is only one element in this partition the equilibrium is completely uninformitive, while if elements of the partition each contain exactly one element the equilibrium is completely informative, i.e. the sender communicates the true state perfectly.

From the neuroscience perspective, we chose this problem because it allows us to explore the basis of deviations from "rational" behavior in social situations away from a moral context. In the past few years there have been a number of imaging studies examining the neural correlates of social emotions such as trust and resentment **???**. These have focussed largely on the *emotional* factors of decision making. This paper focuses more on the constraints neural function may place on social reasoning divorced from moral implications. Cheap talk is a form of strategic information transmission, which means that it may include strategic deception. The ability to deceive is non-trivial from a neural perspective since it requires the deceiver to understand that another person may have a different perception of the world than their own. Children do not usually develop this ability until they are about 4 years old. From the perspective of the person receiving information, responding to possible deception may be difficult since it requires a decision maker to ignore some or all of the information given. Economists almost implicitly assume that all information from another decision maker is suspect until proven otherwise (i.e. through reputation building or shared goals), taking this sort of neural perspective may indicate that the brains default is to trust information.

One well known task in both psychology and neuroscience is the "Stroop Task". In this task subjects are show the names of colors in different colored ink. At the begining of each trial the subjects are told to either read the word or name the color of the ink. There are essentially two types of stimuli: congruent stimuli where the word and color match, eg. red in red ink; and incongruent stimuli where the word and color as missmatched, eg. blue in green ink. MacDonald et al found that the dorsolateral prefrontal cortex was differentially activated by the color-naming task, the more cognitivially difficult task, than the word-reading task implying that this area may be instrumental in implementing cognitive control, whereas they found that the anterior cingulate cortex was differentially activated by incongruent versus congruent stimuli. We find significant correlations to both of these areas during our task, supporting the hypothesis that departures from rational behavior may be the result of automated information processing in the brain.

3. The Bargaining Task

Subjects are recruited in pairs and read detailed instructions. Following experimental economics conventions, all the details of the protocol (except the specific values of the buyer in each stage) are known by both players. This convention means that it is possible for players to arrive at an economic equilibrium, in theory, simply by reasoning and introspection (so if they deviate from equilibrium, it is not because they were confused or misled). It also means that noise from perceptions about the nature of the experiment and the information of the other player are minimized, so we can study neural activity that arises solely from perceptions of the other players' hidden value and strategy.

At the beginning of each round the players were told whether they were the "buyer" or the "seller" in that round. Buyers were told their random private value for a hypothetical good, distributed identically and independently across trials over integers 1 through 10. Sellers never had any intrinsic value for the good (as if they were producing perishable goods at zero marginal cost), to simplify the game. In this simple structure, trades should always take place since the buyer's value is always above the seller's cost of zero (i.e. trade is *always* efficient). Subjects switched roles every five rounds.

The task had two stages at which subjects make choices affecting the game:

- Stage 1: The buyer is informed of her private value, v and asked to "suggest a price". We call the buyer's suggested price s.
- Stage 2: The seller sees s and is asked to choose a price p. If $p \leq v$ the trade occurred in that round; the buyer receives the difference between the value and the price (the surplus in economic terms) $v - p$ and the seller receives the net profit p.

FIGURE 1. Screens seen by the buyer and seller during one trial.

In addition we asked subjects directly about their beliefs. For the seller this was always a first-order belief: what do you think the buyer's value is. However for the buyer this is always a second-order belief, i.e. a belief about someone else's beliefs: what does the seller *think* your value is. Subjects were rewarded for their answers to these questions according to their accuracy, getting one point for the correct answer and a decreasing payoff as their answers diverged from the correct answer. This way, we incentivized honest reporting, but the payoffs for belief elicitation were far outweighed by the payoffs to the game itself.

Because we wanted to image brain activity of both the buyer and seller in a pair simultaneously, we used a partner protocol in which two subjects play each other repeatedly. Games like this are sensitive, in theory, to what two partners know about the history. This reputational history-dependence is a huge source of complication in analyzing brain activity since the signal-to-noise ratio is so low in fMRI. It is therefore usually necessary to pool many trials which are informationally similar, making feedback impractical. Therefore neither player received any direct feedback about the outcome of each trial. Subjects were not even told whether a trade had taken place. This allows us to treat each trial as approximately independent, and to analyze behavior as if they are participating in a one-shot game.

Each subject performed this task 60 times, 30 as a buyer and 30 as a seller, switching roles every 5 rounds. The role switching allows us to compare behavior as buyer and seller within each subject, which is statistically advantageous. They were not informed of their total earnings until the very end of the experiment. Since the sellers never received information about the buyer value, subjects never learned directly about their opponent's suggested pricing strategy. Similarly, since the buyers never learned whether a trade occurred, they could not learn directly about the seller's pricing strategy. However, sellers could conceivably make inferences based on the *range* of suggestions they received over course of the experiment, a possibility we investigate below.

A theoretical analysis of the one shot game shows that if a seller's strategy is at all sensitive to suggested prices, high-value buyers will all pool with the low-value buyers, transmiting minimal information and making the seller's strategy non-optimal (see appendix for formal details). So the equilibrium prediction is that suggested price should be completely unrelated to value ($Prob(v|s) = .1$ for all v, s), and the seller will choose p maximize $u(p)(11 - p)/10$, for a risk neutral seller this implies that he should always choose $p = 5$ or $p = 6$.

Keep in mind that in the theory above, subjects are forming *correct* beliefs about what other players will do and maximizing their own personal gains given their beliefs. As a useful benchmark contrast, suppose both players were trying hard to cooperate and earn the most money together from the task (E.g., suppose they had planned beforehand to split their earnings equally, so they want to maximize the total gain.) Since the seller always wants to sell the good to maximize their joint gains, the seller should try to always name a price the buyer can afford. One way to do this is to always name a price of 1. Another way to do this is for the buyer to make a suggestion which is communicates the value perfectly, i.e. the suggestion correspondence has a functional inverse $f(s) = v$, and for the seller to choose $p \leq f(s)$. Such patterns could emerge, in theory, but neither emerged in this experiment, instead we observe partial information revelation of the sort you might expect if the buyer and seller incentives were somewhat aligned.

4. Behavioral Results

In a study like this which is aimed at two audiences, it is always helpful if the behavioral patterns themselves are surprising and of interest to readers who are not interested in neural activity. Roughly speaking, buyers' suggested prices are often remarkably revealing of their values. Buyers not only sent suggestions which are statistically informative (in the sense that s is realted to v), but they seemed to use linear strategies resulting in bins similar to those predicted by Crawford and Sobel for games where player's incentives are partially aligned. Furthermore, sellers seem to anticipate informativeness of price suggestsions and set their prices p based on the suggestions s also using roughly linear strategies. Interestingly, however, a single player's strategies in their role as buyer were not generally best responses to their own strategies as a seller (and vice versa).

Figure 1 reports all the data from a single subject. The top graph shows the subject's suggested prices s against her valuations v. The bottom graph shows the prices chosen by that same subject, p against the suggested prices she saw s. Both look roughly linear and have much more statistical association than predicted by equiliibrium theory. As a buyer, she seems to "shave" prices by suggesting a fraction of the value. As a seller, she seems to follow a markup strategy of adding a small constant to the buyer's suggestion.

Looking at these behavioral data for each subject suggests that they can be characterized by three parameters in each role: The slope and intercept of the regression line of suggested prices s on values v (as buyer) or prices p on suggestions s (as seller), and the associated R^2 of each regression. The buyer-regression slope is a measure of information revelation while he seller-regression slope is a measure of information-sensitivity.

We can express all three parameters compactly for both roles by plotting each subject's buyer and seller regression parameters in a two-dimensional "intercept-slope" space, in the figure below, with intercepts on the x-axis and slopes on the y-axis. Since each subject has two sets of parameters (because they played both buyer and seller roles) their parameters are connected by a chord. Red circles represent seller regressions (p on s) and blue circles represent buyer regressions (s on v). The sizes of the circles for each subject role are proportional to the R^2 statistic for the regression. The particular subject whose full suggested price and price regressions are graphed in figure 1 is the one with the filled circles.

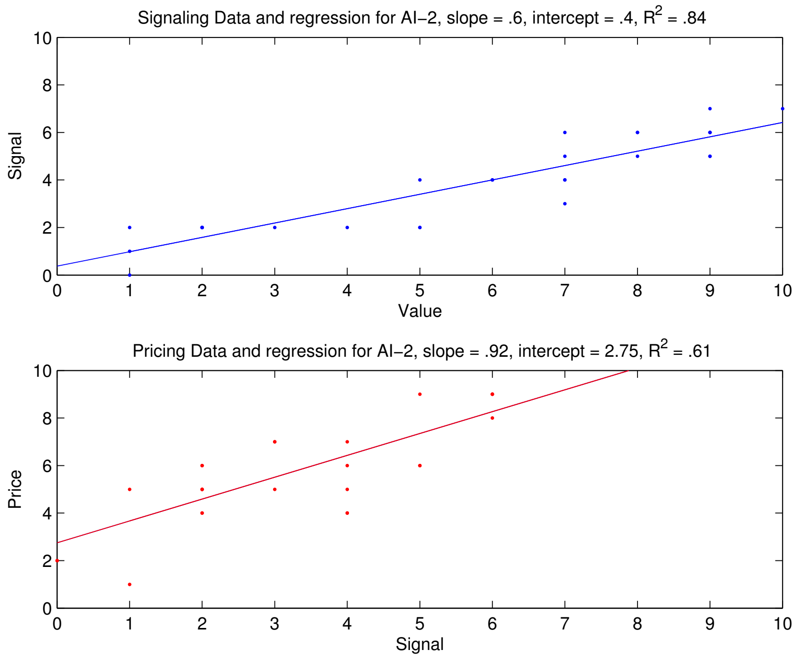

Figure 2. Suggested prices sent by a single subject as buyer, and prices set as seller. (Points are jittered by adding random noise so that identical points are plotted separately).

To make the graphs more readable, subjects were divided into four groups based on the R^2 statistics for the two regressions, which represent the behavioral predictability of their strategies. The majority of subjects have substantial predictability in both s and p (R^2¿.30 for both regressions). The most common pattern is like that shown in Figure 1, the red (seller) circle has a positive intercept and slope around 1, which indicates the sellers are marking up the suggested price by a fixed constant, while as buyers, the blue circles typically have slopes around .5 and small intercepts. Note that the presence of a significant markup by the seller indicates that these subjects are not trying to be "fair" by suggesting $s = \frac{v}{2}$, since that would imply that they should also price fairly by choosing $p = s$.

The four subjects shown in the graph on the lower right show little information revelation as buyers (the theorized "babblers"), but as sellers, they still seem to mark up suggested prices. For the 5 subjects in the upper and lower left graphs, nothing systematic is happening[1]. Recall that the simple theory prediction, babbling and ignoring the suggested price, predicts that players' parameter estimates will be centered around slopes of zero, i.e. no information revelation or information sensitivity. In addition it predicts that the seller

[1]However, as one might expect the 3 subjects got some of the lowest payoffs

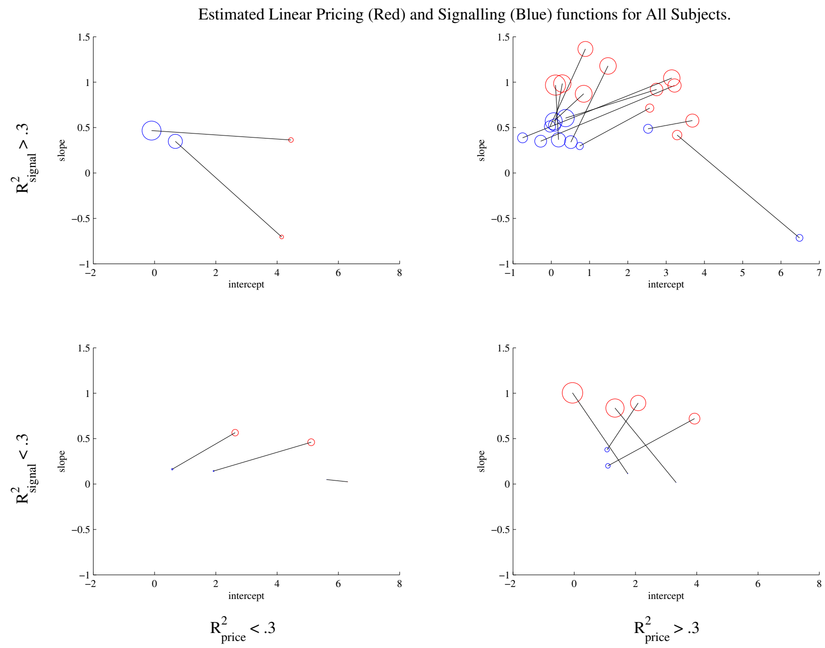

Figure 3. Behavioral data for all 20 subjects separated into quadrants according the R^2 statistics for the p on s, R_p^2, and s on v, R_s^2, regressions.

R^2 statistics should be close to one, since they should employ a constant strategy. The theory clearly does not fit well in general, although there are many cases where slopes are low we see nothing corresponding to the predicted seller strategies since these all have very low R^2 statistics. Furthermore, there is striking cross-subject variation which might be explained by imaging.

The extra information transfer by the subjects resulted in an increased level of efficiency. Equilibrium predicts that $50-60\%$ of the feasable trades should take place. We found that 68% of trades took place over all the 20 buyer-seller pairs; in eight of the 20 pairs at least 75% of the feasible trades were made. One pair traded in every round, extracting all of the possible surplus. The efficiency of pairs was weakly correlated with the information revelation and information sensitivity of the buyer/seller pair.

5. fMRI results

The central brain imaging dependent variable is the blood oxygen level-dependent (BOLD) signal. This is a coarse measure of brain activity based on how much oxygenated blood flows through a brain area with a short time lag (2-4 seconds). The generic analysis correlates the time series of the BOLD signal in each artificially-defined voxel (a 3.4x3.4x4 mm cube) with a matched time series of some sort of regressor.

Generally we look for clusters of $k \geq 5$ contiguous voxels which each have BOLD activity correlated with the regressor at a low p-value (typically $p < .001$). There are several types of analyses one can do with BOLD-signal data. We will present only our results derived so far which illustrate how these analyses work and what we have learned and might be learned with further analysis.

5.1. **Between-individual analysis.** A between-individual analysis takes some behaviorally derived parameter for each person (or, perhaps each pair) and regresses the brain activity linked to some task, such as the onset of a stimulus, on this parameter. Any such correlation permits a statement of the form "higher activity in area A is linked to the propensity to exhibit behavior B". These analyses are quite difficult because brain activity is measured with a time lag and error, and individual-level behavioral parameters are often measured with error as well.

There are two interesting brain areas whose activity at the time of price choice is correlated with the seller's information sensitivity (the slope in Figure 1) in a between-individual analysis. Activity in the anterior cingulate (ACC) is *decreasing* in slope and activity in the right superior temporal sulcus (RSTS) is *increasing* in slope (Figures 4 and 5 show sections and scatterplot).

The cingulate is typically active during cognitive conflict or "executive function". A class of tasks which reliably produce cingulate activity are called "Stroop tasks" (after the pioneering psychologist Colin Stroop). Stroop gave subjects words printed in ink colors which sometimes matched the word and sometimes did not. The task is then to rapidly name the ink color, but not the word. When subjects see the word they often mistakenly say the word rather than name the ink color. The child's game "Simon says" is also a Stroop task; for an American in England, looking to the right side for oncoming traffic (rather than the familiar left side) is a Stroop task too. More generally, a Stroop task requires a decision maker to override a rapid, highly practiced, automatic response. The fact that ACC activity is higher when sellers are less sensitive to the buyer's price suggestions hints that ignoring a suggested price is like overriding an automatic response in a Stroop task.

The STS region, whose activity is positively correlated with seller's information sensitivity, is known to be active in "theory of mind" tasks, in which people must form judgments about what others intend, it is particularly associated with detecting biological motion, or listening to meaningful speech **?**. The STS

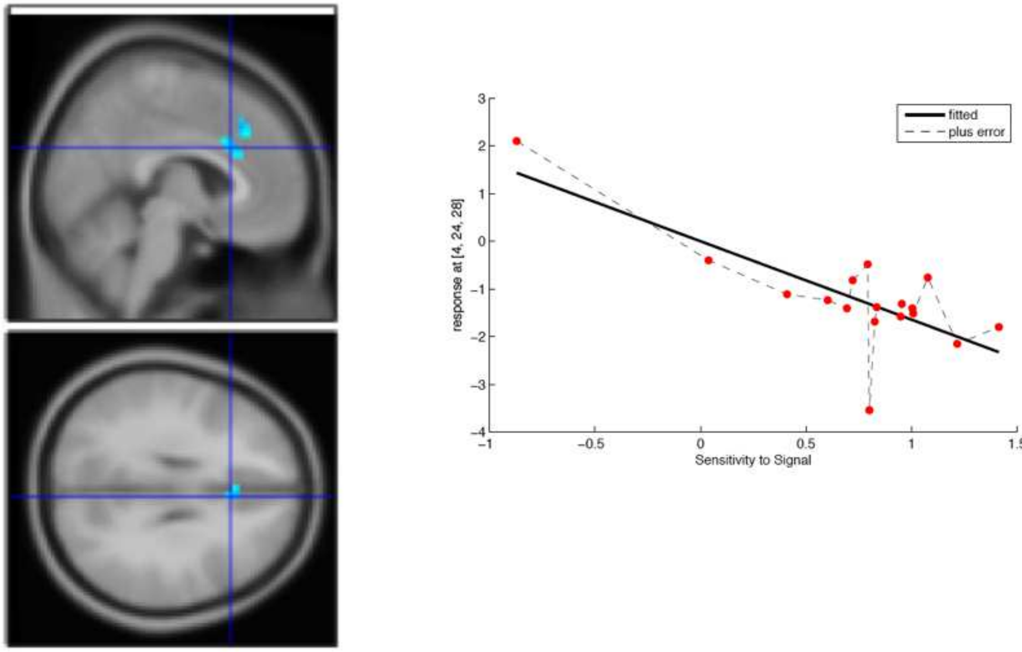

FIGURE 4. Activity in the anterior cingulate cortex at the point of price choice is negatively correlated with a subject's information sensitivity. Activation is shown overlain on the MNI average brain on the left, the plot of relative activation against information sensitivity is shown on the right.

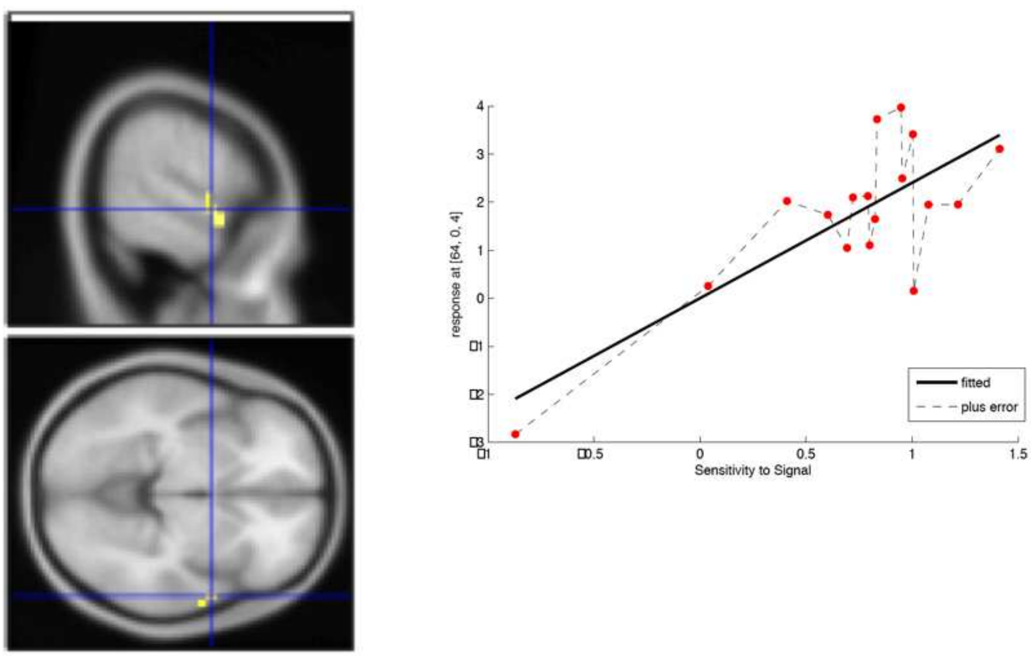

FIGURE 5. Activity in the right superior temporal sulcus is positively correlated with subject information sensitivity. Activation is shown on the left, parameters are plotted on the rights.

and ACC activity together suggest two patterns of behavior different subjects adopt: In one pattern, sellers with high information sensitivity do not perceive a conflict between the suggested price s and the optimal price p (since parametrically, the relation between those two variables is high), so ACC is relatively inactive; instead, the STS is active since these sellers are attributing meaning to these suggested prices. Conversely sellers with low information sensitivity exhibit little STS activity because they are not attributing a lot of meaning to the suggested price. However, because the price they choose and the suggested price are not correlated, the ACC is busy trying to resolve the conflict between the "automatic" response to the suggested price and the "override" of choosing a different price.

5.2. **Parametric analysis.** An emerging type of analysis searches for areas in which the BOLD signal is correlated with a parameter of behavior on a trial-by-trial basis. We have performed two such analyses which yield interesting results.

First, when the buyers are considering what price to suggest, there is a positive correlation between suggesting low prices and activity in the left STS (i.e., a negative correlation between price and LSTS activity) (see Figure 6). The left STS is involved in both producing and interpreting language. Activity here indicates that when buyers are planning to name low prices they are using brain regions also used in thinking about what these signals mean, and how others might interpret messages. Notice that this activation is espcially high for *low* suggested prices, i.e. trials where the buyer is likely trying harder to effect the seller's beliefs. Interestingly, this correlated prain activity is locked to the the onset of the decision task rather than the actual choice, i.e. these signals may actually be able to predict the buyer's choice of suggested price.

Second, when sellers are choosing prices, the "markup" they choose over the suggested price $(p - s)$ is correlated with activity in right dorsolateral prefrontal cortex (RDLPFC) (see Figure 7) This area, like the ACC, is involved in cognitive control and working memory. In fact, as noted above, the left DLPFC is differentially active during the Stroop task when subjects are asked to name the ink color instead of read the word. It is also activated when responders in ultimatum games receive low offers (compared to high offers; Sanfey et al 2004), and when players in games make choices rather than express beliefs about choices

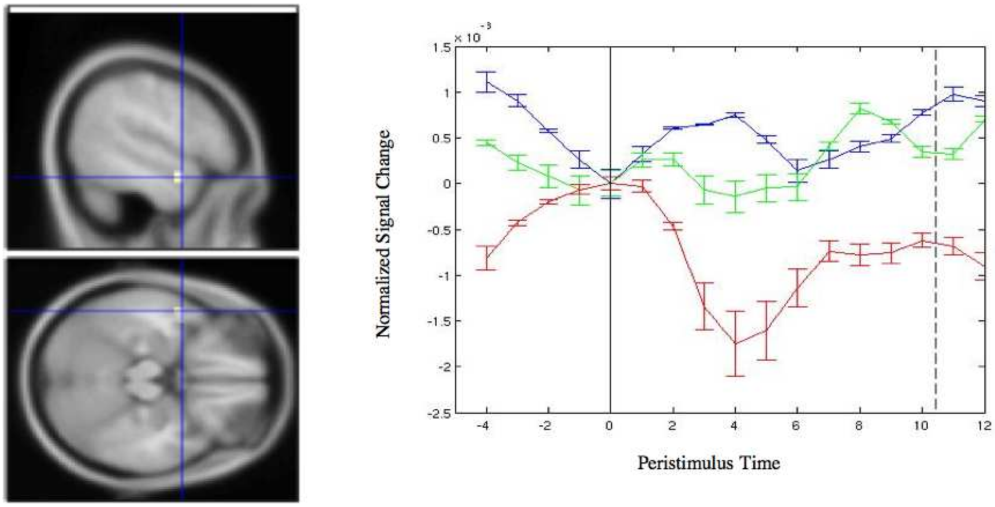

Figure 6. The left STS is correlated within subject to suggested price choice. The cluster is shown on the left, peristimulus activation plots are shown the right for s chosen in the lowest (blue), middle (green) and highest(red) thirds of each subjects choice range. The solid vertical line is the time of stimulus onset, the dotted vertical line is the average decision time.

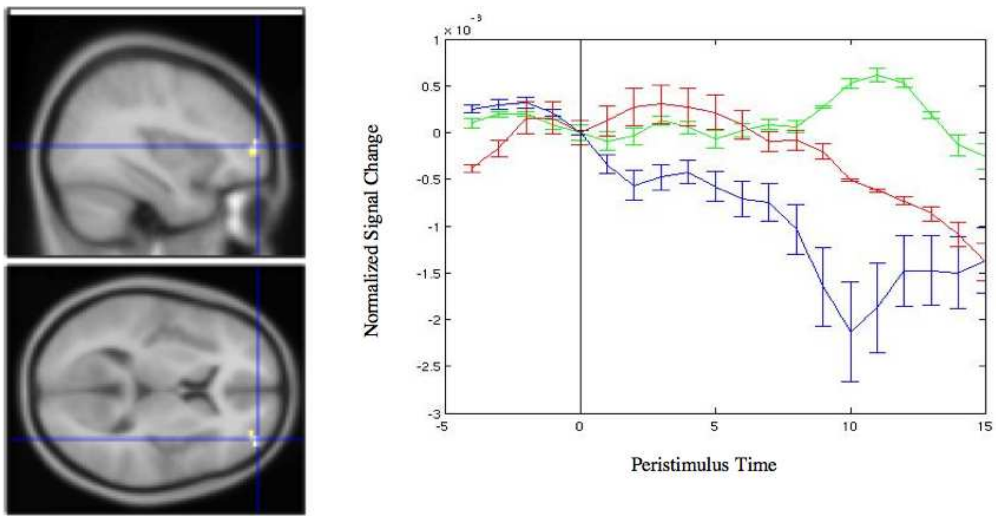

Figure 7. The right DLPFC is correlated withing subject to the sellers chosen markup and the price decision point. Once again, the cluster is shown on the left on the MNI average brain with peristimulus activation plots for trials divided into the top (red), middle (green) and low (blue) thirds of each subjects range.

of others (Bhatt and Camerer, GEB 2004). When sellers choose larger markups, they are exerting more cognitive control to override the impulse to use the information they receive as is, i.e. the suggested price.

5.3. **Predicting behavior from neural activity and future experiments.** One way neural measures might be especially useful for economics is if they can add predictive power to other types of variables.

Neuroscientists refer to the attempt to predict behavior from neural signals, which precede the behavior in time, as reverse inference. We plan to include some analyses of this type in a future draft of this paper.

In our study, a candidate prediction of a causal change comes from the well-studied ACC. It is known that scarce cingulate resources can be tied up if subjects have to perform a parallel activity which requires attention or working memory (e.g. remembering a several-digit number). Suppose sellers also have to remember a number while they are deciding what price to choose. If ACC resources are substituted away from the pricing decision and toward number-memory, then the negative correlation between ACC activity and information sensitivity suggests sellers will become *more* information sensitive (i.e., they will respond with prices more closely linked to). That is, adding a number-memory treatment could change the outcome of the bargaining. Another candidate prediction comes from the positive correlation between right DLPFC activity and higher price markups (seller prices above buyer-suggested prices). Knoch et al (2006) successfully used TMS to disrupt brain activity in right DLPFC and induce players in ultimatum games to accept low offers more frequently. Their interpretation is that the DLPFC is activated by a conflict between the desire to accept money and the desire to enforce social norms by rejecting a low offer. TMS disruption disables this region so behavior reverts to the simpler, more natural reaction of acceptance. Similarly, one prediction from the DLPFC-markup link we see is that if TMS was used to disrupt activity in the right DLPFC in our task, markups would be reduced.

We don't know if either of these experiments would have the predicted effect, or whether they are worth doing at all. We offer these suggestions, at this exploratory point, mostly as illustrations of the recipe of using the fMRI to identify regions of activity which can then be influenced by other treatments, in a way that potentially changes behavior.

6. Conclusion

The behavioral data from this experiment show that subjects have a strong tendency toward both information revelation and information sensitivity, even in situations where economic theory predicts there should be none. The neural data linked to these behaviors suggest that rather than a social or emotional explanation (subjects want to be fair), this information transfer may be the result of a far more basic automated response to information. In fact, some activations seen in this study are very close to those seen in the classical information processing experiment, the Stroop task. These sorts of data suggest that we should reconsider what the "default" response of a decision maker might be to any piece of information.

Appendix A. Bayes-Nash equilibria of the one-shot game

Claim A.1. *The only Bayes-Nash equilibrium of the one-shot game is one where the buyer's suggestion is uninformative.*

Intuitively this follows from the fact that in this game the buyer and seller's incentives are almost exactly opposed. They both prefer that a trade take place, but given that a trade does take place the game is constant sum. To prove this formally we first need to define a few functions.

The probability mass functions for both player's strategy based on the information they receive. For buyers this is the perfect information v, while for sellers this is the noisy and possibly uninformative suggestion s.

$$\begin{aligned} s(x|v) &= Pr\{suggestion = x|v\} \\ p(x|s) &= Pr\{price = x|s\} \\ q(v|x) &= \frac{s(x|v)}{\sum_y s(x|y)} \end{aligned}$$

The utility functions for the buyer and seller respectively are:

$$\begin{aligned} U^B(s,v) &= \sum_{x<v} p(x,s)u_b(v-x) \\ U^S(p,s) &= u_s(p)\sum_{v\geq p} q(v|s) \end{aligned}$$

where u_s, u_b and increasing, weakly concave functions on $\mathbb{R}$ such that $u_b(0) = u_s(0) = 0$.

Let n be the lowest price chosen by the seller with positive probability over all suggestions and choose s^* such that $p(n|s^*) > 0$. Let S_0 be the set of suggestions such that $p(n|s) > 0$. We will show that any buyer with value $v > n$ will always choose a suggested price in S_0 by induction. First note that since n is the lowest price charged given any suggested price, suggestions in S_0 strictly dominate other strategies for buyers with value $n+1$ since any other price yields a payoff of 0.

Assume that all buyers with values in $[n+1, m-1]$ all only choose suggested prices in S_0. So any $s' \notin S_0$ $q(v|s') = 0$ for all values between $n+1$ and $m-1$, which in turn implies that $Pr(v \geq p|s') = Pr(v \geq m)$ for all $v \in [n+1, m-1]$. In other words, given s', m dominates all prices in $[n+1, m-1]$. By assumption $p(x, s') = 0$ for all $x \leq n$ so we have $p(x, s') = 0$ for all $x < m$. Therefore $U^B(s', m) = 0$ for any $s' \notin S_0$ and buyers with value m will always choose $s \in S_0$. Note that this means that all buyers with value $v < n$ must always choose $s \in S_0$ as well since otherwise there would be some $s \notin S_0$ such that $Pr(v < n|s) = 1$ implying that $y(p|s) > 0$ for some $p < n$.

Lemma A.1. *In equilibrium sellers choose prices according to some fixed distribution $p^*(x)$ regardless of the suggested price received.*

We show this using an induction argument like the one above. We need to show that $p(x|s_1) = p(x|s_2) = p^*(x)$ for all $s_1, s_2 \in S_0$ that are sent with positive probability. To see this note that it must be true for $x = n$ since if there is any s_1, s_2 such that $p(n|s_1) > p(n|s_2)$ s_1 dominates s_2 for buyers with value $n+1$. Once again, this implies that $s(s_2|n+1) = 0 \Rightarrow p(n+1|s_2) = 0$ and s_1 dominates s_2 for buyers with value $n+2$. Itterating up we see that s_1 must dominate s_2 for all buyers and $s(s_2|v) = 0$ for all v. Once again, using induction we see that $p(x|s_1) \cong p(x|s_2)$ for all s_1 s_2 sent with positive probability. In other words, any suggested price actually sent by the buyer in equilibrium has no effect on the seller's pricing strategy. Let S_1 be the set of suggestions sent in equilibrium.

Let x be some price sucht that $p^*(x) > 0$. This means that x is a best response to $s^* \in S_1$, specifically it must be at least as good as $x+1$, and $x-1$ so

$$\begin{aligned}
u_s(x)\sum_{v\geq x} q(v|s^*) &\geq u_s(x+1)\sum_{v\geq x+1} q(v|s^*) \Rightarrow \\
u_s(x)q(x|s^*) &\geq (u_s(x+1) - u_s(x))\sum_{v\geq x+1} q(v|s^*) \Rightarrow \\
u_s(x)s(s^*|x) &\geq (u_s(x+1) - u_s(x))\sum_{v\geq x+1} s(s^*|v)
\end{aligned}$$

Since $p^*(x)$ is the same for all $s^* \in S_1$ we can sum over all the s in S_1 and get

$$u_s(x) \sum_{s^* \in S_1} s(s^*|x) \geq (u_s(x+1) - u_s(x)) \sum_{v \geq x+1} \sum_{s^* \in S_1} s(s^*|x).$$

We know that buyers will always choose $s^* \in S_1$, so this is simplified to

$$\frac{u_s(x)}{u_s(x+1) - u_s(x)} \geq 10 - x.$$

Since u_2 in increasing and concave the left side of this inequality is strictly increasing while the right side is strictly decreasing. Also notice that if the inequality is *strict* $p^*(x+1) = 0$, so the seller will mix between at most 2 prices x and $x+1$ (this occuring only when the intersection point x^*, such that $10 - x^* = \frac{u_s(x^*)}{u_s(x^*+1) - u_s(x^*)}$ is an integer). This in turn implies that the seller simply chooses x to maximize $(11 - x)u_s(x)$ the expected utility when the seller has no information.