Exploring the Nature of "Trading Intuition"*

Antoine J Bruguier[1], Steven R Quartz[1] and Peter L Bossaerts[1,2]

[1]California Institute of Technology, Pasadena, CA 91125, USA

[2]EPFL, CH-1015 Lausanne, Switzerland

This version: 28 July 2008

Abstract

The Efficient Markets Hypothesis (EMH) and the Rational Expectations Equilibrium (REE) both assume that uninformed traders can correctly infer information from the trading process. Experimental evidence starting in the mid 80s has consistently confirmed the ability of uninformed traders, even novices, to do so. Here, we hypothesize that this ability relates to an innate human skill, namely, Theory of Mind (ToM). With ToM, humans are capable of discerning malicious or benevolent intent in their environment. We confirm our hypothesis in evidence that participation in markets with insiders engages the (evolutionary relative new) brain regions that specialize functionally in ToM, and in evidence that skill in predicting price changes when there are insiders is correlated with scores on two traditional ToM tests. Since ToM is generally understood to rely on pattern recognition (which recently has been confirmed formally in the context of strategic games), we searched for features in the data with which one could identify the presence of insiders. We discovered that GARCH-like persistence in absolute price changes in calendar time characterizes markets with insiders.

I. Introduction

This paper reports results from experiments meant to explore how uninformed traders manage to read information from transaction prices and order flow in financial markets with insiders. Since the seminal experiments of Charles Plott and Shyam Sunder in the early 1980s [Plott and Sunder (1988)], it has been repeatedly confirmed (and we will do so here too) that uniformed traders are quite capable of quickly inferring the signals that informed traders (insiders) receive about future dividends, despite the anonymity of the trading process, despite lack of structural knowledge of the situation, and despite of the absence of long histories of past occurrences of the same situation from which they could have learned the statistical regularities.

It is striking that so little is understood about the ability of the uninformed to infer the signals of others. This ability constitutes the basis of the efficient markets hypothesis EMH [Fama (1991)], which states that prices fully reflect all available information. Underlying EMH is the idea that the uninformed will trade on the signals they manage to infer, and that, through the orders of the uninformed, these signals are effectively amplified in the price formation. In the extreme, prices will fully reflect all available information. Without a better understanding of how uninformed practically manage to read information in prices, EMH remains a hypothesis without a well-understood foundation.

The feedback from trading based on inferred information to price formation has been formalized in the concept of the rational expectations equilibrium (REE) [Green (1973), Radner (1979)]. For economists, REE forms the theoretical justification of EMH. But again, REE takes the ability of the uninformed to correctly read information from prices as given, rather than explaining it. As such, like EMH, REE lacks a well-articulated foundation.

The goal of the experiments that we report on here can be expressed in a more mundane way, as an attempt to better define what is meant with "trading intuition," and to understand why some traders are better than others. Books have been written to elucidate trading intuition [Fenton-O'Creevy, Nicholson, Soane and Willman (2005)] but attempts at formalizing the phenomenon have so far failed.

What distinguishes markets with insiders is the presence of a *winner's curse*: sales are often successful only because prices happen to be too low (relative to the information of the insiders), while purchases may occur only at inflated prices. Either way, the uninformed trader is hurt. While the winner's curse is usually associated with strategic, single-sided auctions, it also applies to competitive, double-sided markets, and indeed, the winners' curse is not only implicit in the theory of REE but also very much of concern in real-world stock markets [Biais, Bossaerts and Spatt (2008)]

From the point of view of the uninformed trader, the winner's curse conjures up an image of potential malevolence in the trading process. Humans are actually uniquely endowed with the capacity to recognize malevolence (as well as benevolence) in their environment, a capacity that psychologists refer to as *Theory of Mind* [Gallagher and Frith (2003)]. Theory of Mind is the capacity to read intention or goal-directness in patterns, through mere observation of eye expression [Baron-Cohen, Jolliffe, Mortimore and Robertson (1997)], or movement of geometric objects [Heider and Simmel (1944)], in the moves of an opponent in strategic play [McCabe, Houser, Ryan, Smith and Trouard (2001), Gallagher, Jack, Roepstorff and Frith (2002), Hampton, Bossaerts and O'Doherty (2008)], or in actions that embarrass others ("faux-pas") [Stone, Baron-Cohen and Knight (1998)]. Theory of Mind is thus *the ability to read benevolence or malevolence in patterns in one's surroundings*. This contrasts with prediction of outcomes generated purely by, say, physical laws, which may harm or be beneficial but without intention.[1]

Perhaps the most telling example of how Theory of Mind concerns detection of intentionality in patterns concerns a case where simple physical laws of motion are violated. Subjects are first shown an object (say, a block) attracted to a target (ball) but encounters a physical barrier (a wall) that it overcomes by moving around it (see Figure 1, top). Subsequently, the barrier is re-moved and two situations are shown. In the first one, the block, freed of the obstacle, moves along a straight path towards the original target location (Figure 1, middle); in the second one, the block continues along the original trajectory (Figure 1, bottom). The first situation accords with standard physics; the second one is unusual (from a physics point of view) and may alert the subject that the original path around the barrier was intentional, i.e., that it reflected goal-directed behavior. It has been repeatedly observed that humans (including one-year old infants [Gergely, Nadasdy, Csibra and Biro (1995)] and some non-human primates [Uller and Nichols (2000)] spend *more* time gazing at the second situation than the first one, despite its novelty (the ball takes a different path), suggesting that the suspected intentionality draws their attention.

The first goal of our study was, through experiments, to determine to what extent trading intuition and ToM are related. ToM, as mentioned before, is the ability to read benevolence or malevolence in patterns in one's environment. So, it was reasonable to conjecture that it applied to order flows in financial markets with insiders as well. Specifically, we asked whether the uninformed engage ToM brain regions when shown the order flow of a market with insiders, and whether ToM skills correlate with success in forecasting market prices when there are insiders.

The answers to both questions are affirmative. We found that inside information engaged (only) functionally specific ToM regions in the brain of the uninformed, and we found that ability to correctly forecast price changes in markets with insiders correlated significantly with performance in standard ToM tasks.

The second goal of our study was to start formalizing the concept of ToM to the extent that it applied to financial markets with insiders. ToM is about pattern recognition – which is intuitively clear from the example in Figure 1, and was recently confirmed formally for play in strategic games [Hampton, Bossaerts and O'Doherty (2008)]. For markets, however, it is not even known whether there are any features at all in the data that would allow one to merely identify the presence of insiders (in analogy with the violation of physical laws in the bottom panel of Figure 1 that prompts many to conclude that there may be intentionality). Existence of such features is an important component when arguing that ToM may apply.

Here, we report results from a statistical investigation of trade flows in our experimental financial markets through which the presence of insiders could be discriminated. We find that GARCH-like features emerge when there are insiders. Specifically, autocorrelation coefficients of absolute price changes in calendar time are significantly more sizeable.

It seems to us that the link between GARCH-like features and insider trading is a major finding that deserves further investigation irrespective of its import for the plausibility of ToM in the context of financial markets.

Remainder is organized as follows. Section II describes the experiments and discusses the results. Section III presents the findings from statistical contrasts of trade flows in sessions in our experiments when there are insiders and when there are none. Further discussion is provided in Section IV.

II. Description Of The Experiments And The Results

Here, we provide brief explanations of the experimental set-up (a full discussion can be found in the Appendix) and we present the main empirical results.

We first ran a markets experiment for the sole purpose of generating order and trade flow in a controlled setting, to be used in the main experiments. The latter consisted of: (i) a brain imaging experiment, meant to discern how subjects attempted to interpret the data by localizing areas of the brain that are active during re-play of the markets; (ii) a behavioral experiment, where we tested for subjects' ability to predict transaction prices, and to ascertain their generic ToM skills (we also ran a test of logical and mathematical thinking, for a reason to be discussed later).

II. A. Step 1: Trading Data Collection Experiment

Twenty (20) subjects participated in a parameter-controlled market experiment that used an anonymous, electronic exchange platform (jMarkets; see http://jmarkets.ssel.caltech.edu/). The following situation was replicated several times; each replication will be referred to as a session. Subjects were endowed with notes, cash, and two risky securities, all of which expired at the end of a session. The two risky securities ("stocks") paid complementary dividends between 0 and 50¢: if the first security, called stock X, paid *x* cents, then the second security, called stock Z, would pay 50-*x* cents. The notes always paid 50¢. Allocation of the securities and cash varied across subjects, but the total supplies of the risky securities were equal; hence, there was no aggregate risk, and, theoretically as well as based on observations in prior experiments [Bossaerts, Plott and Zame (2007)], prices should converge to levels that equal expected payoffs; that is, risk-neutral pricing should arise. Subjects could trade their holdings for cash in an anonymous, continuous open-book exchange system. Subjects were not allowed to trade security Z, however. Consequently, risk averse subjects that held more of X than of Z would want to sell at risk-neutral prices; the presence of an equal number of subjects with more of Z than of X allowed markets to clear, in principle. In most of the sessions, to be referred to as *test sessions*, a number of subjects (the "insiders") were given an estimate of the dividend in the form of a (common) signal within $0.10 of the actual dividend. All subjects were informed when there were insiders; only the insiders knew how many insiders there were. Subjects were paid in cash according to their performance and made $55 on average. More details can be found in the appendix.

Figure 2 displays the evolution of transaction prices throughout the experiment. Trading was brisk, independent of the type of session; on average, traders entered or cancelled an offer every 0.7s and one transaction took place every 3.2s. During test sessions, the uneven distribution of information evidently skewed the transaction prices. Overall, the evolution of prices is consistent with prior experiments [which did not control for aggregate risk, however; Plott and Sunder (1988)].

As mentioned before, the sole purpose of the markets experiment was to generate order flow and transaction data from markets with tight control over endowments, information and exchange platform. The resulting data formed the input for the brain imaging and behavioral experiments. For this reason, we will not elaborate on the markets experiment. We will come back to relevant features of the trading flow in the next section, however.

II. B. Step 2: fRMI Experiment

Eighteen (18) new subjects were shown a replay of the 13 previously recorded market sessions in random order while being scanned with fMRI. These subjects played the role of outsiders who did not trade. Subjects were given the instructions of the markets experiment, so that they were equally informed as the outsiders in that experiment. Subjects first chose whether they would bet on stock X or Z, after being told whether there were insiders in the upcoming session or not. Subsequently, the order flow and transaction history of stock X was replayed in a visually intuitive way (see Figure 3 and Video 1). During replay, subjects were asked to push a button each time they saw a trade, indicated by a 500ms (millisecond) change in color of the circle corresponding to the best bid or ask.

Our design had three main advantages. First, the subjects did not trade during the replay (they did have to take positions before the replay). While the question of how the human brain makes financial decisions is interesting, we needed first to understand how humans perceive a stock market. By not introducing decision-making, we avoided a confounding factor. Second, the periods without insiders were perfect controls. Since the trading data acquisition method, the display screens, and the number of traders were the same, the two types of sessions were identical in every respect except for the presence or absence of insiders. Third, by adding a blind bet, we elicited a feeling of “randomness.” Indeed, if we had forced subjects to choose stock X for every session, the payoff would have been the same fixed number for every subject. Moreover, we could not have separated an increase in stock price from a higher expected reward, as these two signals would have been perfectly correlated. Instead, by introducing a blind bet, we orthogonalized expected reward and stock price.

To identify areas of significant brain activation in the insider sessions relative to the control sessions, we used a standard approach as implemented in the package BrainVoyager (Brain Innovation, Maastricht, The Netherlands). We fit a General Linear Model (GLM) to the (filtered, motion-corrected) activation data for each "voxel" (a cubic volume element of $27mm^3$) which included several auxiliary predictors (to capture motion and visual effects) as well as the expected reward for the subject and four task-relevant predictors (Figure 4). The task-relevant predictors were as follows. First, we constructed a parametric regressor that would quantify the effect of the insiders (if present) on stock prices. We used the absolute value of the difference between the stock's trading price and 25¢. Separate parametric predictors were thus computed for the sessions with insiders and without insiders. This way, the sessions without insiders could be used as controls, against which the sessions with insiders could be compared. Second, we added block regressors (dummy variables), to identify sessions with and without insiders.

All predictors were convolved with the standard "hemodynamic response function," in order to account for the lag between activation at the neuronal level and the fMRI signal. Areas of significant brain activation "contrast" were then recovered as those areas where the difference (across insider and no-insider sessions) in the slope coefficients of the predictors were deemed to be significant using standard statistical tests.

Our way of modeling brain activation had several advantages. First, as mentioned above, we stripped out the effect of decision-making in the brain and focused on the perception of the market. Second, the predictors we built were orthogonal to expected reward as we used the absolute value of the distance to 25¢ rather than the (signed) difference with 25¢. Third, by contrasting sessions with and without insiders and making sure that these two types of sessions displayed no obvious differences in trading activity, we controlled for differences in visual activation and isolated the effect of insiders on subjects' perception of the market.

We found a significant contrast for the parametric regressors in one large region of paracingulate cortex (PCC; Figure 5-a and Table I). We also found significant activations in a smaller region in the frontal part of the anterior cingulate cortex (-14; 23; 39; 5 voxels; Figure 5-a and Table I). Finally, we found strong differential activations in right amygdala (21; -10; -12; 5 voxels; Figure 5-b and Table I) and left insula (-30; -7; 11; 5 voxels; Figure 5-c and Table I).

Significant contrasts for the block predictors showed up in a large area of lingual gyrus (-9; -65; -6; 25 voxels; Figure 6 and Table II), as well as a small area of cerebellum (-13; -58; -30; 9 voxels; Figure 6 and Table II). No other areas with five or more voxels exhibited significant contrasts.

Our results provide support for the hypothesis that Theory of Mind is involved when subjects are facing markets with insiders. We found that the same brain areas are activated as during traditional ToM tasks. PCC (paracingulate cortex) activation is standard in tasks involving ToM [Gallagher and Frith (2003)], and in strategic games in particular [Gallagher, Jack, Roepstorff and Frith (2002)], [McCabe, Houser, Ryan, Smith and Trouard (2001)]. Similar activation was also reported in choice vs. belief game theory experiments [Bhatt and Camerer (2005)]. The PCC has also been observed in tasks that involve attribution of mental states to dynamic visual images, such as intentionally moving shapes [Castelli, Happe, Frith and Frith (2000)].

We also found that activation in the right amygdala and the left anterior insula increased as transaction prices deviated from the uninformed payoff. While some studies have reported the involvement of these structures in ToM tasks [Baron-Cohen, Ring, Wheelwright and Bullmore (1999), Critchley, Mathias and Dolan (2001)], they are more typically regarded as involved in processing affective features of social interaction. Specifically, the amygdala is a critical structure in the recognition of facial emotional expressions of others [Phillips, Young, Scott, Calder, Andrew, Giampietro, Williams, Bullmore, Brammer and Gray (1999), Phan, Wager, Taylor and Liberzon (2002), Morris, Ohman and Dolan (1998)] while the anterior insula is thought to play a critical role both in subjective emotional experience [Bechara and Damasio (2005)] and in the perception/empathetic response to the emotional state of others [Singer, Seymour, O'Doherty, Kaube, Dolan and Frith (2004)]. A complementary interpretation of these activations is that they may also reflect subjects' own emotional responses to market activity, which is consistent with recent psycho-physiological evidence that financial market participation engages somatic marker (emotional) circuitry during heightened market volatility [Lo and Repin (2002)]. Future research should shed more light on this potential link between market participation and emotions.

We also found activations in the frontal part of the ACC (anterior cingulate cortex). Our activation is in a slightly more posterior and dorsal location than when ToM is used in strategic and non-anonymous, simple two-person games [Gallagher, Jack, Roepstorff and Frith (2002), McCabe, Houser, Ryan, Smith and Trouard (2001)].

The increased activation of lingual gyrus in the presence of insiders provides further support for the role of ToM in market perception. For example, the lingual gyrus is involved in the perception of biological motion, a key cue for mentalizing [Servos, Osu, Santi and Kawato (2002)]. However, increased activation of lingual gyrus may also be related to recent accounts that this structure is involved in complex visual tasks where subjects are asked to extract global meaning despite local distractors [Fink, Halligan, Marshall, Frith, Frackowiak and Dolan (1996), Fink, Halligan, Marshall, Frith, Frackowiak and Dolan (1997)]. When there are no insiders, subjects can concentrate on the task we imposed, namely, to track trades. In our display, transactions were a local feature, indicated by changes in color of circles in the middle of the screen. In contrast, when there are insiders, the entire list of orders may reflect information with which to re-evaluate the likely payoff on the securities but at the same time, subjects are still asked to report all transactions, which now amounts to a local distraction. Our interpretation of lingual gyrus activation is inconsistent with theories [Easley and O'Hara (1992), Glosten and Milgrom (1985)] that predict that only the trade prices are relevant to update beliefs. If these theories were right, other features of the display need not be watched, and simultaneously keeping track of transactions is not a distraction. Rather, our finding is consistent with experimental evidence that insiders do not exploit their superior information by exclusively submitting market orders [Barner, Feri and Plott (2005)]. Future research should determine to what extent lingual gyrus activation reflects ToM (through motion of objects) or the proverbial conflict between the "forest" (insider information) and the "trees" (trades).

We did not observe any significant activation in brain areas related to formal mathematical reasoning. In particular, there was no evidence of estimation of probabilities [Parsons and Osherson (2001)] or arithmetic computation [Dehaene, Spelke, Pinel, Stanescu and Tsivkin (1999)]. We also did not observe any significant activation in brain areas related more generally to problem-solving or analytical thought [Newman, Carpenter, Varma and Just (2003)] or reasoning [Acuna, Eliassen, Donoghue and Sanes (2002)]. This finding has also been highlighted in a recent study of ToM in strategic games [Coricelli and Nagel (2008)]. We shall elaborate in the next Section.

It is important to note that engagement of ToM brain structures was not simply in response to a complex graphical representation of the market, as both control and test sessions generated approximately equivalent market activity. Similarly, in both conditions the market activity was the result of human interactions. Hence, it was not merely the presence of human activity that resulted in ToM, as this was equivalent in both conditions. Rather, the salient difference between session type was the fact that, during sessions with insiders, trading activity provided information that could be extracted with the appropriate inferences.

II.C. Step 3: Behavioral Experiment

The behavioral experiment was meant to correlate subjects' ability to predict trade prices in markets with insiders and their general ToM skill as reflected in traditional ToM tests. We added a test for mathematical and logical thinking, prompted by the absence of activation in regions of the brain usually associated with formal mathematical and probabilistic reasoning in our task.

Forty-three (43) new subjects were given a series of four tasks that were administered in random order. The first task was a Financial Market Prediction (FMP) task in which the trading activity was replayed to the subjects at original speed and paused every 5 seconds. During half of the pauses, we asked subjects to predict whether the next trade was going to occur at a higher, lower, or identical price as the previous trade that had happened. For the other half of the pauses, we reminded subjects of their predictions and informed them of their success. The second task was a ToM task based on eye gaze [Baron-Cohen, Jolliffe, Mortimore and Robertson (1997)]. The third task was a ToM task based on displays of geometric shapes whose movement imitated social interaction [Heider and Simmel (1944)]. As with the FMP task, we paused the movie every five seconds and asked subjects to predict whether two of the shapes would get closer, farther, or stay at the same distance. For the other half of the pauses, we reported the outcomes and the subjects' predictions. The fourth task (the M task) involved a number of standard mathematics and probability theory questions [and frequently used in Wall Street job interviews –see Crack (2004)].

The behavioral experiment confirmed the use of ToM. We observed a significant correlation (p=0.048 and p=0.023) between the FMP task and both ToM tasks (Figure 7, top panels) but no significant correlation (p=0.22) between the FMP and the M tasks (Figure 7, bottom left). Surprisingly, we did not observe any correlation between the two ToM tests (Figure 7, bottom right). Self-reports during the behavioral experiment did not show any evidence of personalization. We did not observe any significant gender differences.

Our behavioral experiments corroborate the relationship between trading success and ToM that our fMRI results suggest. The correlation between the two ToM tests and the ability to predict price changes in a financial market confirms that successful traders use ToM. We also found that the subject scores on the two ToM tests were not correlated (Figure 7-D). This finding suggests that ToM is a multi-dimensional set of abilities, a possibility that has not been raised yet in the ToM literature: each test represents one dimension of ToM abilities, but financial abilities correlate with both dimensions.

Lack of (significant) correlation between the scores on the two ToM tests provides indirect confirmation that general intelligence or state of attentiveness cannot explain the significant correlations between scores on the FMP and ToM tests.

IV. Theory of Mind in Markets with Insiders: What Patterns To Attend to?

Formally, ToM remains a rather elusive concept. It is often defined only vaguely, and mostly in terms of specific tasks [Gallagher and Frith (2003)], or in terms of activation of particular brain regions [McCabe, Houser, Ryan, Smith and Trouard (2001), Gallagher, Jack, Roepstorff and Frith (2002)]. It is generally accepted, however, that ToM concerns pattern recognition, but only in simple examples is it immediately obvious what patterns are involved. Figure 1 (bottom panel) provides an example: the motion of an object violates physical laws and hence is readily recognized as revealing intention.

In the context of strategic games, however, one recent study has successfully identified the patterns in the moves of one's opponent on which ToM builds. Thus, in collaboration with Alan Hampton and John O'Doherty [Hampton, Bossaerts and O'Doherty (2008)], one of us established that, when a subject is engaged in playing the inspection game, brain activation in ToM regions of the human brain reflect the encoding of an error in predicting *how one's opponent changes moves as a result of one's own actions*.[2]

The import of this finding is that ToM in game play can be concretized in terms of precise mathematical quantities that characterize an opponent's actual play. As such, ToM concerns concrete, "online" learning of game play. This is consistent with the proposition that ToM involves pattern recognition – detection of the nature of intentionality in one's environment.

It deserves emphasis that ToM therefore contrasts with Nash reasoning, where players would simply hypothesize that opponents choose Nash equilibrium strategies and that they would stick to them. Unlike ToM, Nash reasoning is "offline:" it works even if one never sees any move of one's opponent; it is also abstract: it concerns what the opponent could rationally do and how to optimally respond. Consistent with the idea that ToM and Nash thinking are not related, brain regions that are known to be engaged in abstract logic and formal mathematics [Dehaene, Spelke, Pinel, Stanescu and Tsivkin (1999), Newman, Carpenter, Varma and Just (2003), Acuna, Eliassen, Donoghue and Sanes (2002)], do not display significant activation during game play [Coricelli and Nagel (2008)].

In the context of our markets with insiders, it is by far clear on which patterns ToM could build. This is because it has not even been established whether there are any patterns that distinguish markets with and without insiders. But our finding that subjects engage in ToM to comprehend insider trading and the generally accepted view that ToM concerns pattern recognition suggests that they exist.

Consequently, we set out to determine whether there are features of the data that differentiated markets with insiders. In the physical movement of objects, violation of physical laws signals the presence of intentionality (see Figure 1). In markets, we hoped to identify statistical properties of the data that are not present when there are no insiders. We focused on an examination of the trade flow.

We looked at a host of time series properties of the trade flows in the markets experiment that formed the basis of our study, such as duration between trades or skewness in transaction price changes. In the end, it was persistence in the size of transaction price changes over calendar time intervals that provided the only statistically significant discrimination. As such, GARCH-like features appear to distinguish markets with and without insiders.

Specifically, we computed transaction price changes over intervals of 2s (as mentioned before, there is a trade every 3.7s on average, so calendar time tick size was chosen to be slightly shorter than average duration between trades). We followed standard practice and took the last traded price in each 2s interval as the new price, and if there was no trade during an interval, we used the transaction price from the previous interval.

We then computed the first five autocorrelations in the size (absolute value)[3] of transaction price changes. Figure 8-a plots the results for two adjacent periods, Periods 7 and 8. As can be inferred from Figure 2, there were 14 insiders (out of 20 subjects) in Period 7, while there were none in Period 8. The patterns in the autocorrelations of the absolute price changes in the two periods are very different. There is substantial autocorrelation at all lags for Period 7 (when there are insiders) while there is none for Period 8 (when there are no insiders). Figure 8-b shows that this is a general phenomenon. Plotted is the sum of the absolute values of the first five autocorrelation coefficients against the number of insiders. We refer to the former as "GARCH intensity" because it measures the extent to which there is persistence in the size of price changes. GARCH intensity increases with the number of insiders; the slope is significant at the 5% level and the R-squared, at 0.32, is reasonably high given the noise in the data.

Consequently, it appears that GARCH-like features in transaction price changes provide one way to recognize the presence of insiders, and hence, a foundation on which ToM in markets with insiders could build. As such, we have shown that there indeed exist patterns in the data that reflect intentionality, a generally accepted condition for ToM to apply. From this perspective, our findings that participation in markets with insiders engages ToM brain regions and that ToM skills and price forecasting performance are correlated make sense.

We have not identified yet what aspect of these GARCH-like features subjects are exploiting to infer inside information, but our results open up many promising avenues for future research. Our approach of combining markets experiments with brain imaging and individual behavioral testing may thus eventually lead to a formal understanding of how humans manage to comprehend a social institution as complex as a financial market.

V. Further Discussion

Our imaging results for markets with insiders have at least one more aspect in common with recent neurobiological studies of ToM in strategic games, namely, the absence of activation in regions of the brain involved in formal mathematical reasoning. The results suggest that abstract reasoning skills are unrelated to performance in either domain. With respect to strategic games, it is not known whether performance is correlated with mathematical skill. With respect to markets with insiders, we set out to investigate the issue. Specifically, as part of the third, behavioral experiment, we tested our subjects on their ability to solve mathematics and logic problems and correlated their scores with performance on the market prediction test. We focused on a set of formal problems that have become popular in recruitment in the financial industry [Crack (2004)]. Figure 6 (Bottom Left Panel) shows that, while the correlation was positive, it was insignificant. The behavioral results are therefore in line with the imaging results.

Our study is not only relevant to finance. It contains at least two major contributions for psychology and neuroscience. We are the first to report activation in ToM brain regions during the perception of anonymous multi-agent systems, significantly extending the scope of ToM. Moreover, along with Hampton, Bossaerts and O'Doherty (2008) and Coricelli and Nagel (2008), we are the first to quantify ToM. In particular, when there were insiders, the more transaction prices deviated from the uninformed payoff of 25¢, the more evidence outsiders had that there was important information to be inferred, making mentalizing increasingly important.

Personalization (or anthropomorphization) is often suggested as an explanation why ToM sometimes applies even in situations that do not directly involve personal contact with another human being, such as in the Heider movie used in one of our ToM tests [Heider and Simmel (1944)]. In the context of financial markets, personalization emerges at times in speech, when markets are described as being “exuberant,” “jittery,” or “anxious” [Oberlechner (2004)]. Answers on the exit survey following our behavioral experiment, however, do not reveal any evidence of personalization.

Hence, it appears that the personal component is not important. Intentionality that can be inferred from patters in one’s environment *is*, however. This should explain why ToM also applies to large-scale anonymous social interaction (such as our financial markets).

The importance of intentionality also clarifies recent evidence [McCabe, Houser, Ryan, Smith and Trouard (2001), Gallagher, Jack, Roepstorff and Frith (2002)] that ToM brain regions are activated when playing strategic games with humans and not when playing against a pre-programmed computer. The evidence should be interpreted *not* as indicating that the human component *per se* was important, *but* as suggesting that potential malevolence was crucial. Indeed, the computers were pre-programmed in a way that was completely transparent to subjects, and did not involve any intention to harm or exploit; they followed specific "laws" just like the object in Figure 1, middle panel, obeyed physical laws.

Our findings have immediate implications for hiring practice in finance. They suggest that the financial industry should test candidate traders on social skills, especially those involving ToM, such as the ability to recognize intent in other people's eyes, or recognize malevolence in the movement of geometric objects, and, because it has recently been shown to also involve ToM, skill in playing strategic games. Our advice is far more specific than the one to come out of a recent analysis of over one hundred professional traders, which revealed that trader compensation (a rather indirect measure of trading success) correlated positively with experience and rank (perhaps not surprisingly) and negatively with psychological measures of emotionality, most of whom are relevant for general everyday problem solving rather than being specific to financial market trading [Fenton-O'Creevy, Nicholson, Soane and Willman (2005)].

Our findings should also help improve visual representation of order and trade flow. Since humans often are best in recognizing the nature of intention in moving (animate or inanimate) objects [Castelli Castelli, Happe, Frith and Frith (2000), Heider and Simmel (1944)], we suggest that traders may be more likely to successfully detect insider trading when order and trade flows are presented in a moving display, as opposed to the purely numerical listings commonly found in the industry. In fact, one may wonder whether the success of our (untrained) subjects in predicting price changes in markets with insiders may be attributable to our using a graphical interface where order and trade flow are translated into movement of circles of various sizes and colors.

In the same vain, one may conjecture that ToM is behind the success of continuous double auctions over one-shot call clearing systems in facilitating equilibration [Plott and Vernon (1978)]. In the former mechanism, the continuous flow of orders may induce ToM, thus enhancing trust that there is no inside information. In a call market, absent order flow information, assessment of the situation cannot rely on ToM and hence must invoke traders' abilities to make the right conjectures about endowments, preferences and information of others. The practical superiority of the continuous auction system over the call markets extends to situations where there are insiders [Plott and Sunder (1988)], even if the latter may in theory be better [Madhavan (1992)].

Appendix

In this appendix, we describe in detail the three experiments of this paper. We used college students or college-educated subjects for every experiment. Both Caltech and UCLA ethics boards approved the experiments.

A. 1. Prior Markets Experiment

We collected trading data with a markets experiment. We asked 20 subjects to trade with the help of an anonymous computerized trading system. They traded for 13 independent sessions. At the end of the experiment, we paid the subjects with cash according to their performance.

Subjects owned two types of stocks. They could buy and sell the first type ("stock X") for a price between 0¢ and 50¢. At the end of each session, this stock paid an amount of money that we drew randomly between 0¢ and 50¢ (uniformly distributed). We called this amount of money "dividend." The subjects learned the value of the dividend only at the end of each session. Subjects could not trade the second type of stock ("stock Z"). It paid a dividend that was complementary to the dividend of stock X (the two dividends added up to 50¢). For example, if we revealed at the end of a session that the dividend of stock X was 38¢, then the dividend of stock Z was 12¢.

For each session, the payment was determined as followed. Before trading started, we endowed each subject with a different mix of stock X, stock Z, and cash. During trading, subjects could exchange stock X for cash as they saw fit. As a result, at the end of trading, a subject generally owned a different combination of stock X, stock Z and cash. After markets closed, we revealed the value of the dividend, and subjects learned the amount of their payoff for the session. For example, if a trader owned 10 units of stock X, 5 units of stock Z, and had 87¢ in cash, when we revealed that the dividend for stock X was 38¢, the trader learned that her payoff was 10×38¢ + 5×(50¢ - 38¢) + 87¢ = $5.27. We added $5.27 to the trader's earnings, and then we restarted a new trading session independently.

Since the subjects did not know the dividend until the trading session was over, the expected payoff of one unit of a stock was 25¢. In addition, risk-averse subjects could neutralize risk by balancing their holdings of the (complementary) stocks X and Z. Since the total endowment of stock X and Z across all subjects was the same, everyone could in theory balance holdings. As a result of this absence of aggregate risk, risk-neutral pricing should obtain. Consequently, without insider information, the predicted equilibrium price is 25¢.

We explained in detail the payoff schedule but did not tell the subjects that the equilibrium price of the stock was 25¢. During the trading, the price moved as predicted by the theory, confirming earlier experiments [Bossaerts, Plott and Zame (2007)]. Essentially, when the price was below 25¢, subjects tended to purchase the stock. Indeed, if a stock that pays off on average 25¢ were trading at, say, 17¢, a subject that would purchase it would earn on average 8¢. The purchase would push the price of the stock higher, until it reached the equilibrium, at 25¢. Similarly, traders would sell a stock trading at, say, 33¢ until the price went down to 25¢.

The setup described above constituted the two control sessions of our experiments. For the 11 other sessions (test sessions), we introduced an additional factor.

For these test sessions, we split the traders into two randomly chosen groups: the "insiders" and the "outsiders." At the beginning of each session, we gave additional information to the insiders in the form of a "signal." The signal was a number chosen at random within 10¢ of the actual dividend of X (uniformly distributed). For example, if we told the insiders that the signal was at 14¢, they knew for sure that at the end of the trading session we would reveal a dividend between 4¢ and 24¢. We did not give any information to the outsiders except the fact that there were insiders in the market.

The uneven distribution of information skewed the trading. At the beginning of each session, the two groups attached different values to the stocks. For example, if we told the insiders that the signal was 14¢, then they would value stock X at

14¢. The outsiders did not have this additional information and could only estimate the value of stock X to be 25¢. Insiders and outsiders had to act carefully. Insiders could use their information to make additional profits. For example, if the signal was at 14¢, they could sell a stock to an outsider who was willing to pay 25¢. The insider would have received 25¢ for a stock that would have paid at most 24¢, and on average 14¢ (i.e. he made an average profit of 11¢). However, by selling, insiders would lower the price of the stock, and thus they would reveal to the outsiders their knowledge of the signal. The outsiders, who knew that there was a signal but did not know its value, had to observe the market carefully and attempt to infer the signal from the trading activity. If the outsiders were successful at estimating the signal from the trade, they would make extra profits. Vice-versa, the insiders had to trade as discreetly as possible to avoid revealing their knowledge of the signal to the outsiders. The anonymity of the trading interface helped conceal their information. At the same time, however, insiders had to trade before the other insiders in order to make additional profits.

The dynamics of such markets are extremely complicated. While researchers have attempted to model them [Admati (1985), Grossman and Stiglitz (1980)], there is no exact description of how the traders (insiders or outsiders) do behave. Thus, we used fMRI to improve our understanding of such markets. We used the periods without insiders as control and the periods with insiders as test.

A. 2. Stimulus Set for fMRI Experiment

We used 18 new subjects and replayed the order and trade flow while we recorded brain activity. First, we explained to the subjects how we have acquired the data and made sure that they understood the experiment by administering a quiz. We also reminded the subjects that they were not going to trade in the market but that they would only observe the replay of a previously recorded market. Still, they had an interest in paying attention as their payoffs depended on what happened in the market (in a way that we describe below). We also instructed them that the term “insider” did not refer to illegal “insider trading.”

We displayed each of the 13 sessions in a different random order for each subject (Figure 3 and Video 1). Each session began with a “blind bet:” we asked subjects to choose between the stock X and the (complementary) stock Z. After the subjects had made a choice, we replayed the market activity for stock X, regardless of their choice, at double the speed (2 minutes and 30 seconds instead of 5 minutes). Finally, we informed the subjects of how much they earned during that session.

For each session, the payment was determined as follows. When subjects placed a blind bet, we endowed them with 10 units of the stock they chose. After the end of the trading replay, we paid them according to the dividend of the stock they had chosen. For example, if a subject had placed a bet on stock X before the trading replay and if we revealed after the trading replay that the dividend was 45¢, we added 10×45¢=$4.50 to the subject’s earnings. If the subject had chosen to bet on stock Z, he would have received 10×(50¢ - 45¢)=$0.50. Thus, the subjects for the fMRI experiments played the role of outsiders who did not trade. They were outsiders because they did not know the signal and they did not trade because they only saw a replay of trades that had occurred in the past.

The stock price was an indicator of the expected reward for the subjects only in the case of a session with insiders. For example, an increasing price indicated that the insiders may have skewed the market because the signal was high. Thus, if a subject had placed a blind bet on stock X, she could have expected her earnings for the period to be high. If she had placed a blind bet on stock Z, she could have expected her earnings for the period to be low. If there were no insiders, the stock price did not contain any information on the expected earning for the period. We illustrated this computation of the expected reward on the third row of Figure 4.

This design had three main advantages. First, the subjects did not trade. While the question of how the human brain makes financial decisions is interesting, we needed first to understand how humans perceive a stock market. By not introducing decision-making, we avoided a confounding factor. Second, the periods without insiders were perfect controls. Since the trading data acquisition method, the display screens, and the number of traders were the same, the two types of sessions were identical in every respect except for the presence or absence of insiders. Third, by adding a blind bet, we elicited a feeling of "randomness." Indeed, if we had forced subjects to choose stock X for every session, the payoff would have been the same fixed number for every subject. Moreover, we could not have teased apart an increase in stock price with a higher expected reward, as these two signals would have been perfectly correlated. Instead, by introducing a blind bet, we orthogonalized expected reward and stock price.

We replayed the market activity (section (iii) of Figure 3) with a simplified interface (Video 1). We represented the price levels for the offers to buy and sell ("bids" and "asks") with a circle. A number inside each circle indicated the price in cents and the diameter of the circle represented the number of units offered. Blue circles represented the bids and red circles represented the asks. When a trade occurred at a certain price, we turned the corresponding circle green for 500ms. We rearranged the circles dynamically to reflect the changes in prices and trades. The locomotion between sessions with and without insiders did not display any obvious differences. Finally, in order to monitor attention, we asked subjects to press a key every time a trade occurred.

Despite the highly salient and vivid nature of our graphical display of the order and trade flow, it should be underscored that it is not what caused the ToM brain activation that we report. This is because we contrast brain activation in a treatment with and without insiders *using the same interface*. Likewise, the human factor behind the market in itself cannot be the cause of the ToM brain activations, because there is an equal amount of human interaction in the control treatment.

A. 3. Construction of predictors

To analyze the fMRI data, we constructed four predictors (Figure 4). We wanted to use the sessions without insiders (second and fourth columns of Figure 4) as controls and compare them with the sessions with insiders (other columns). We constructed two block predictors (last two rows) and two parametric predictors (fourth and fifth rows).

As described above, we could compute the expected reward of one session by taking into account the stock price, the presence or absence of insiders, and the blind-bet (second row of Figure 4). Indeed, during sessions without insiders, neither the blind bet nor stock prices carried information; the expected reward was the mean value of the dividend, namely 25¢. If insiders were present, then a higher stock price for stock X indicated that the dividend was likely to be higher. Thus, if a subject had placed a blind bet on stock X, she would have expected a higher return. Conversely, if she had placed a bet on stock Z, she would have expected a lower return.

We quantified the insiders' activity using the absolute value of the difference between the trading price and 25¢ (Figure 4, third row). Indeed the more insiders skewed the market, the further the price would go from 25¢. We computed this value and built two parametric predictors. One parametric predictor modeled this effect during sessions with insiders (fourth row). To control for other effects, we built a similar predictor for the sessions without insiders (fifth row). By contrasting these two parametric predictors, we isolated the effect of insiders on subjects' perception of the market. We also created two block predictors. They captured the mean brain activation during the sessions with or without insiders.

This modeling had three advantages. First, as mentioned above, we stripped out the effect of decision-making in the brain and focused on the perception of the market. Second, the predictors we built were orthogonal to expected reward as we used the absolute value of the distance to 25¢ and rather than the distance to 25¢. Similarly, our predictors were orthogonal to risk, measured by the bid-ask spread [Glosten and Milgrom (1985)]. Third, by contrasting sessions with and without insiders and making sure that these two types of sessions displayed no obvious differences in trading activity, we controlled for differences in visual activation as well.

A. 4. Behavioral Experiment

This experiment is necessary to confirm the fMRI results because ToM-related brain areas may activate during non-ToM situations.

The experiment consisted of four tasks that we administered in a random order to 43 new subjects. During the first task, we replayed the trading activity for four sessions with insiders. We used the same interface as for the fMRI experiment except that we played the market at the original speed (5 minutes) and we paused the replay of the market every 5s. During these pauses, we alternated predictions and outcomes. Specifically, for half of the pauses, we asked subjects to predict whether the next trade was going to occur at a higher, lower, or identical price as the previous trade that had happened. For the other half of the pauses, we reminded subjects of their prediction and informed them of the outcome. We rewarded subjects for correct predictions but did not punish them for incorrect ones. However, we only gave subjects 5 seconds to make a prediction and penalized them for indecision. Thus, guessing was always better than not answering.

The second task was a ToM test. Using a computerized interface, we recorded the scores on an eye test [Baron-Cohen, Jolliffe, Mortimore and Robertson (1997)]. While several advanced tests were available, such as the faux-pas test [Stone, Baron-Cohen and Knight (1998)], we decided to use this particular one because it provided a continuous scale of ToM abilities and was difficult enough to reveal differences in abilities between healthy adults. In addition, the faux-pas test may not be appropriate as social codes are culturally dependent. We also time-constrained the eye test, and guessing was better than not answering.

The third task was also a ToM test based on displays of moving geometric shapes [Heider and Simmel (1944)]. We played two "Heider movies" and paused the replay every 5s. For half of the pauses, we asked subjects to predict whether two of the shapes would get closer, farther, or stay at the same distance. For the other half of the pauses, we reported the outcomes and the subjects' predictions. The better the subjects understood the social interactions of the shapes, the better they were at predicting movements.

We used the fourth task to measure mathematical abilities. With a computerized interface, we asked subjects to solve seven mathematical puzzles (Table III) under time constraints and without the use of paper. Subjects typed their answers on the keyboard, and guessing was better than not answering.

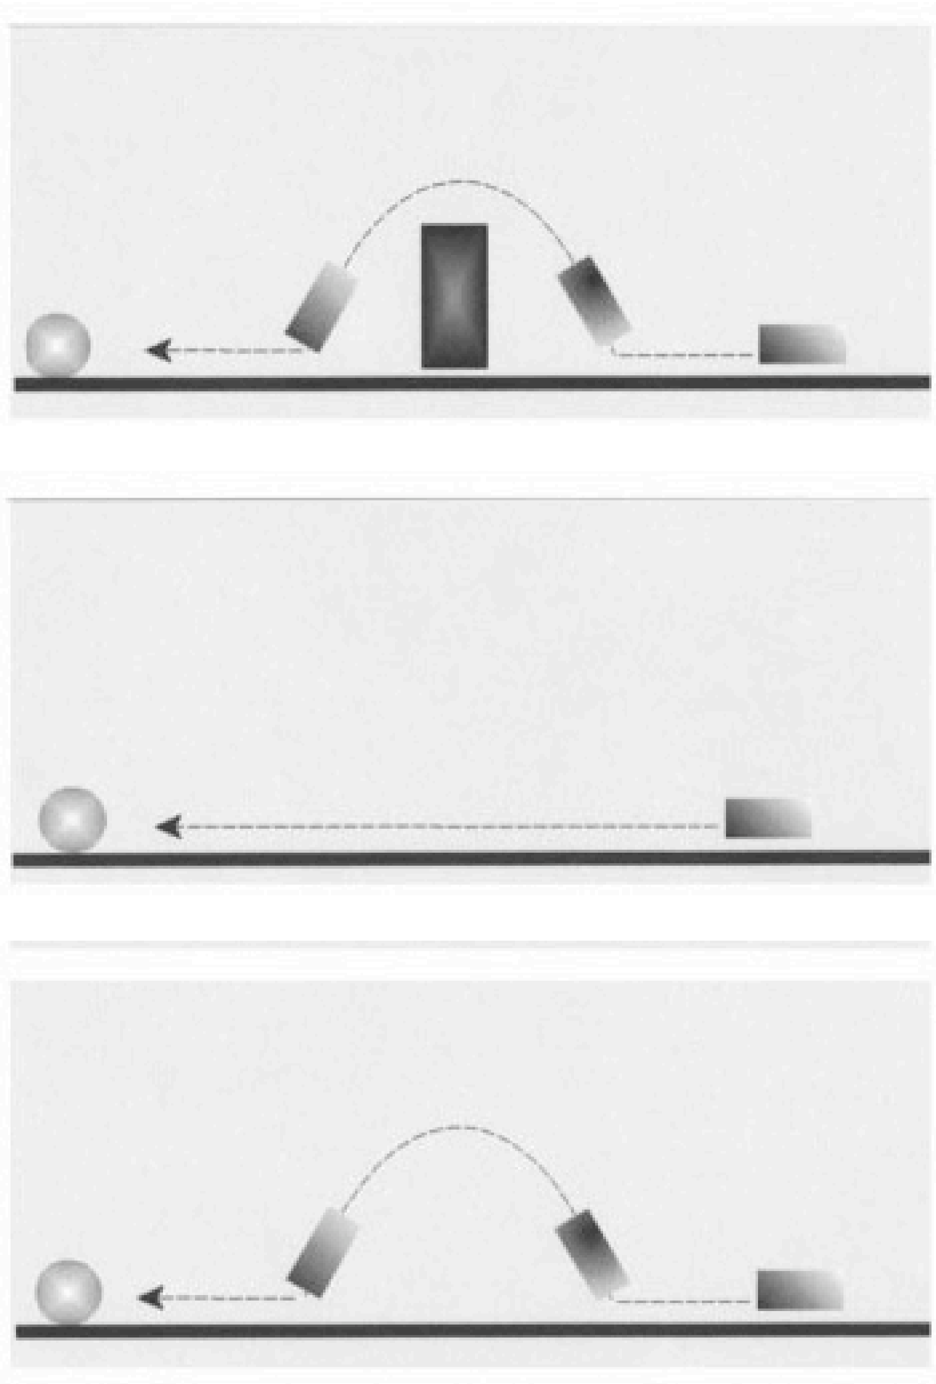

Figure 1. Top: Object (magnetic block at right) attracted to target (ball at left), flowing over wall in the middle; Middle: wall is removed and block moves straight to target, as expected under physical laws; Bottom: Same situation, but now block continues to make a curve even if attracted to target, suggesting *intention* (to curve). (Adapted from: Uller and Nichols (2000).)

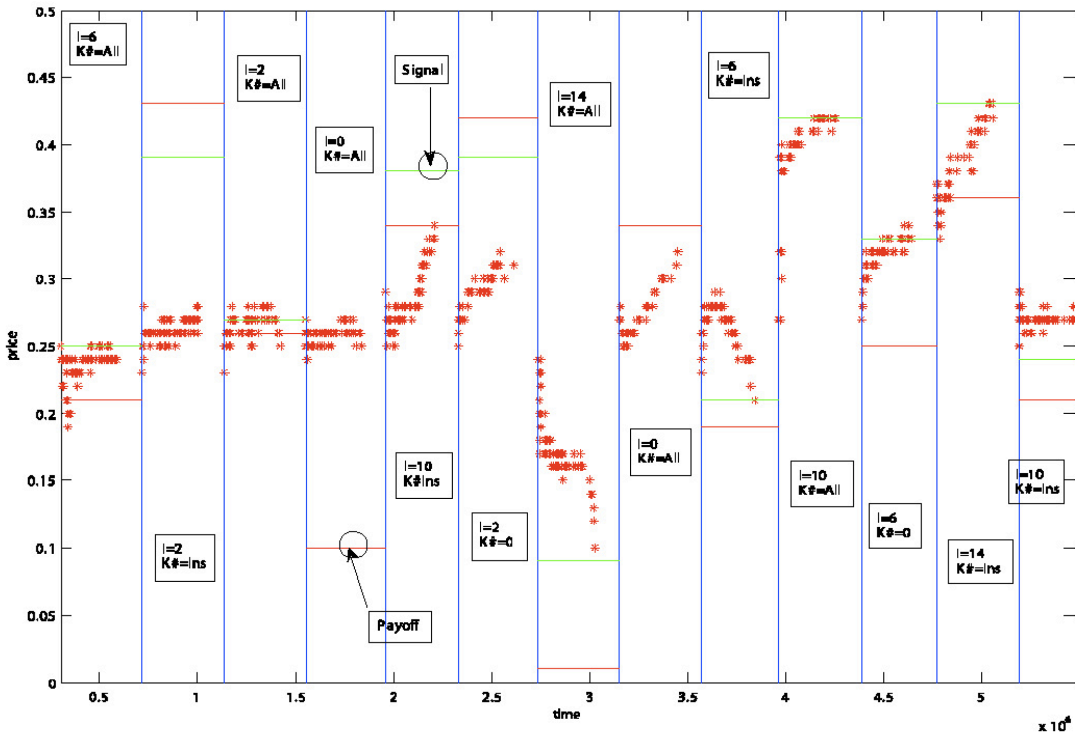

Figure 2. Evolution of transaction prices in the market of stock X. Sessions are delineated with solid vertical separators; dotted vertical separators indicated end-of-trading. Sessions with insiders are marked "ins." Signal levels are indicated with green horizontal line segments; red line segments denote final dividend (of X) for the session. "i" denotes number of insiders (all insiders receive the same signal); "K#" indicates whether all subjects ("All") know how many insiders there were, or none ("0"), or only the insiders ("Ins"). All subjects always knew whether there were insiders.

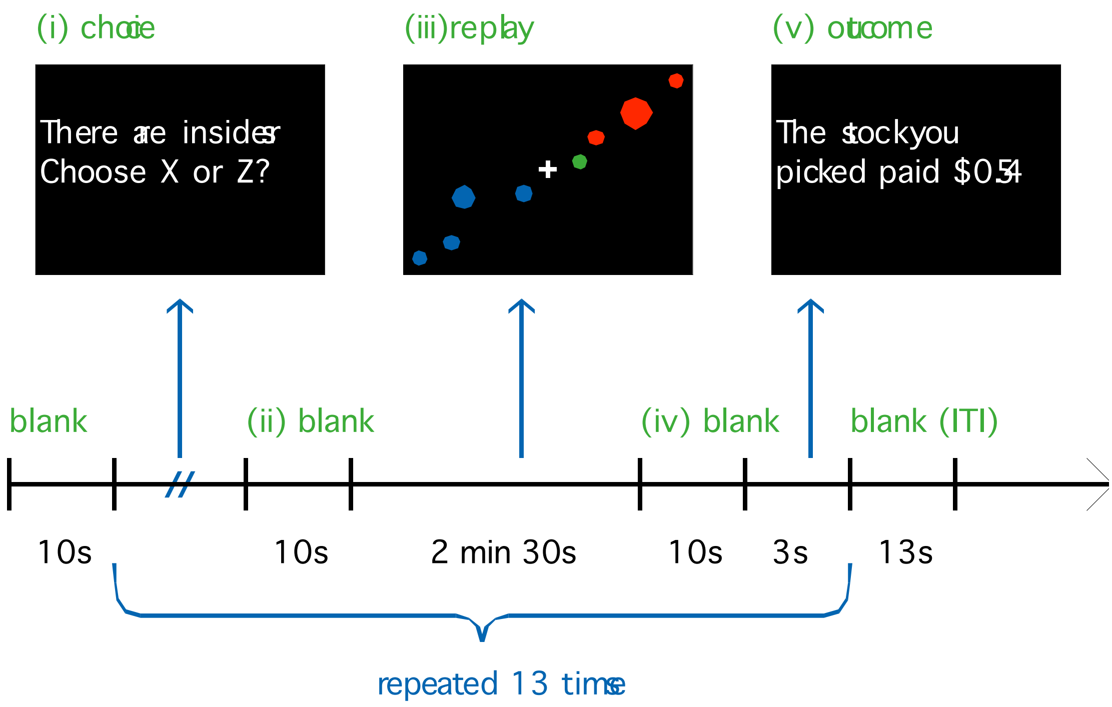

Figure 3. fMRI experiment. (i) At the start of each session, subjects were informed of whether or not the session contained insiders and were instructed to make a choice (a blind bet) between stock X and stock Z. They were then endowed with 10 units of the stock they chose. (ii) A blank screen was then presented for 10 seconds, (iii) a market session was then replayed at double speed, (iv) followed by a 10 second blank screen after which (v) subjects are informed of their payment for that session (10 X their chosen stock's dividend).

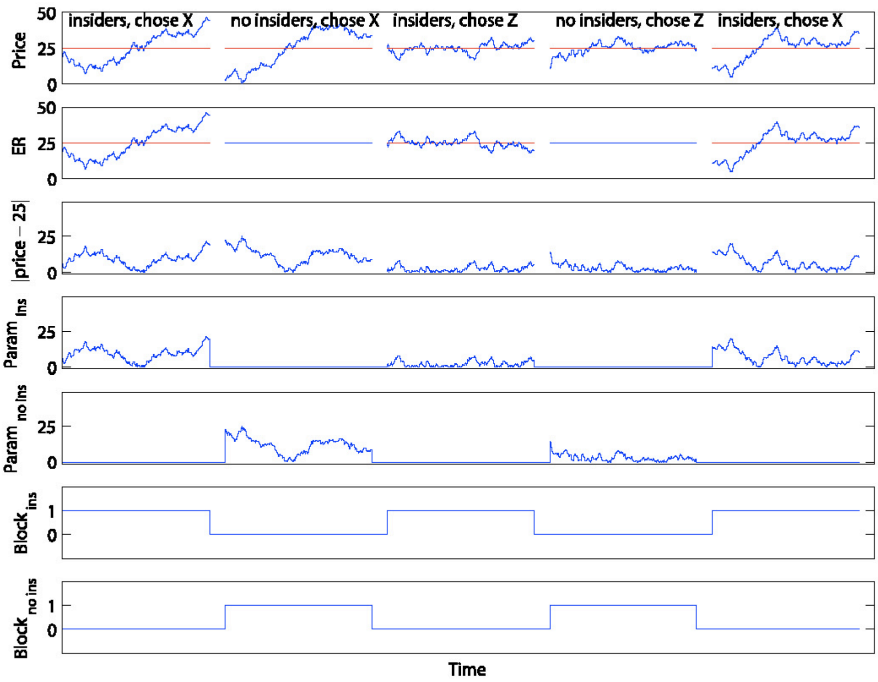

Figure 4. Construction of predictors in the GLM used in the analysis of brain activation data. Columns 1-5 represent five fictive periods of different combinations of presence/absence of insiders and whether the subject chose stock X or Z. A) The evolution of the chosen stock price is displayed. B) The first predictor was the expected reward (ER), computed from the stock price, the presence/absence of insiders, and the blind bet. During sessions without insiders (2nd and 4th sessions), neither the blind bet nor the stock prices carried any information and the expected reward was the mean value of the dividend, namely 25¢. During sessions with insiders (remaining sessions) a higher price for stock X was taken to indicate that the dividend was likely to be higher, resulting in a higher expected reward in the case of a blind bet on stock X and a lower expected reward on stock Z. C) Insider activity was proxied by the absolute value of the difference between the trading price and 25¢. From this proxy, two parametric predictors were constructed. D) First, a parametric predictor modeled this effect during session with insiders. E) Second, a similar predictor was constructed for sessions without insiders. In addition, we added the following predictors: F) Block predictor capturing mean brain activation during sessions with insiders; G) Block predictor capturing mean brain activation during sessions without insiders.

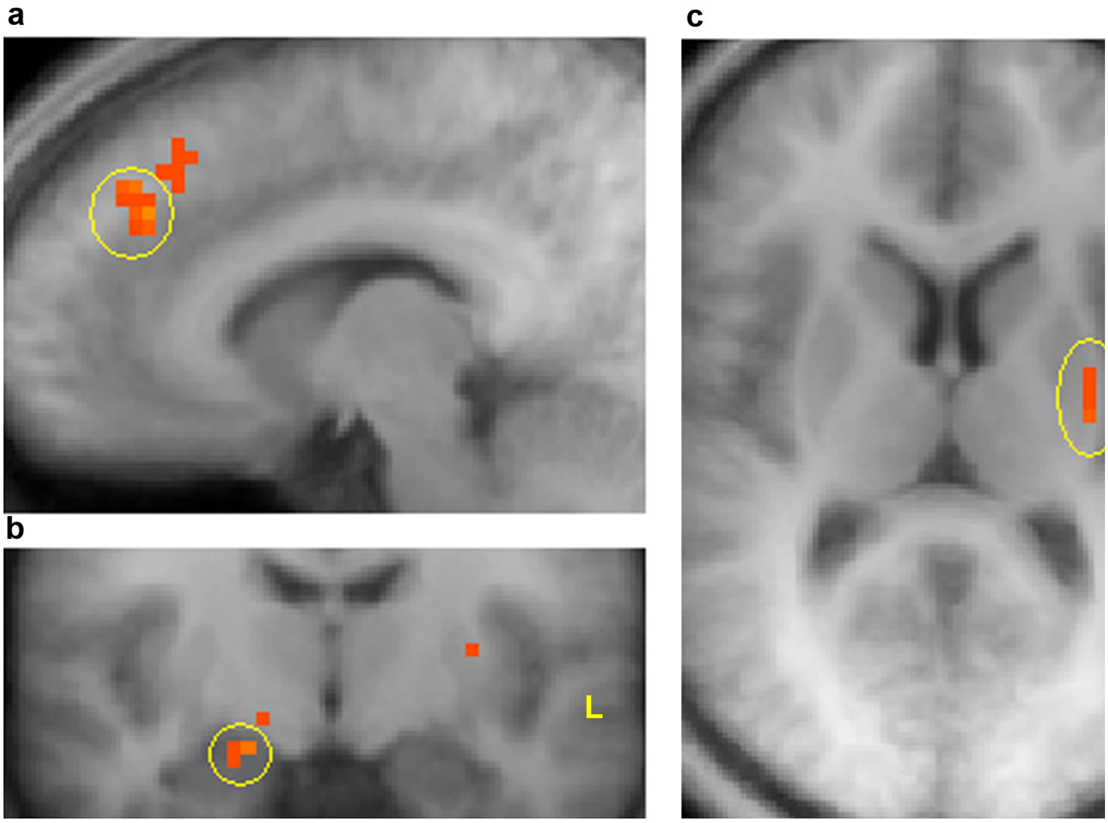

Figure 5. Location of significant contrast between slope coefficients to the parametric regressor between insider and no-insider sessions (p<0.001, uncorrected, random effects). (a) Sagittal view of the activation in paracingulate cortex (Talairach coordinates -9; 41; 36; Brodmann areas 9/32; extends for 22 voxels). (b) Amygdala activation in coronal view (-14; 23; 39; extends for 5 voxels). (c) Activation of the left insula in axial view (-30; -7; 11; extends for 5 voxels).

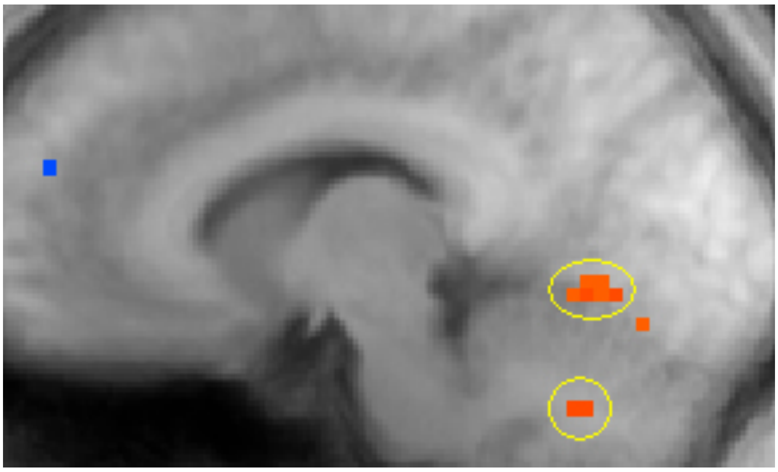

Figure 6. Location of significant contrast between slope coefficients to the block regressor between insider and no-insider sessions. We use the threshold p<0.001 (uncorrected, random effect). We observe two clusters of voxels activated: the lingual gyrus and cerebellum.

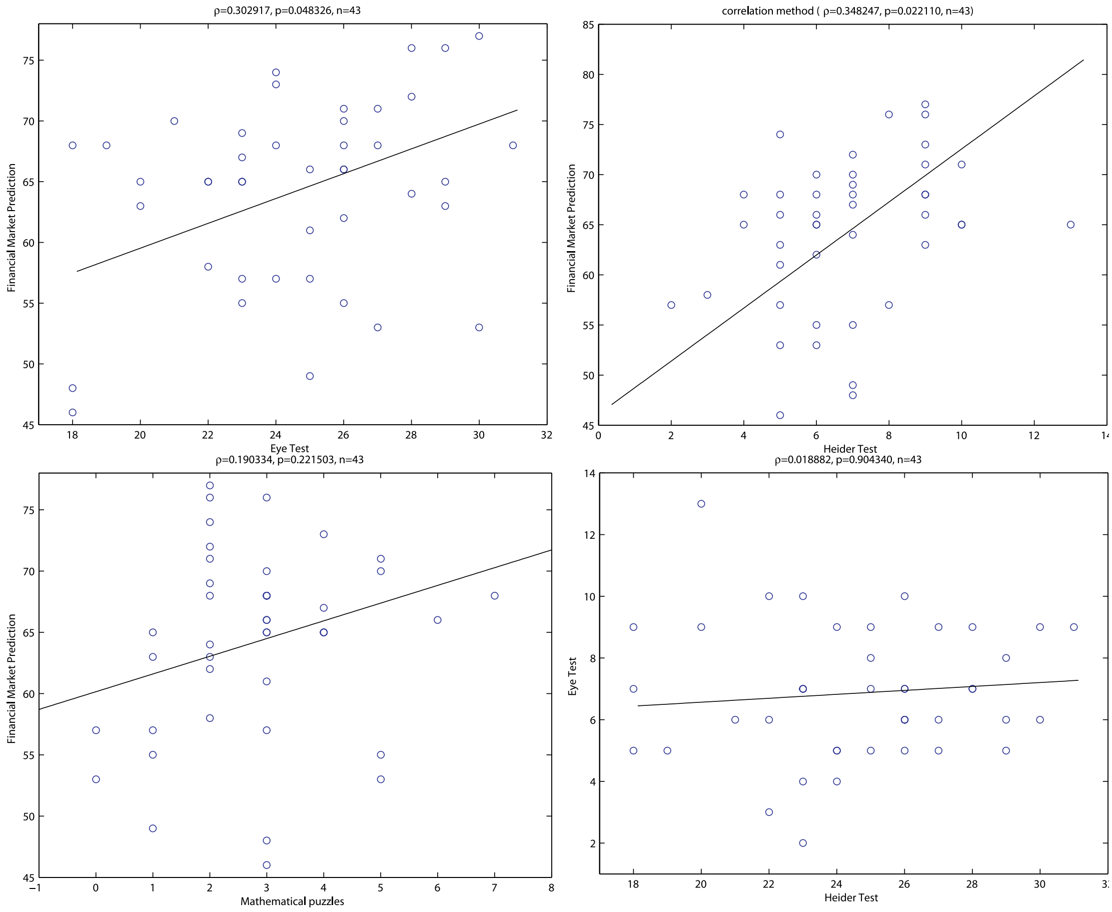

Figure 7. Top panels: Correlation of the scores on the Financial Market Prediction (FMP) and two ToM tests (Left: Eye Test; Right: Heider Test). Since both measures are noisy, we could not use a linear regression and instead we computed the correlation and its p-value. We also computed the mean line with the help of the correlation coefficient. We observe that there is a significant correlation (Eye Test: p=0.048; Heider Test: p=0.023) between the scores on the tests. Bottom Left: Correlation of the scores on the Financial Market Prediction (FMP) and Mathematical (M) tests. While positive, the correlation is not sigificant (p>0.200). Bottom Right: We compare here the performance on the Eye test and the Heider tests. The correlation is insignificant.

a

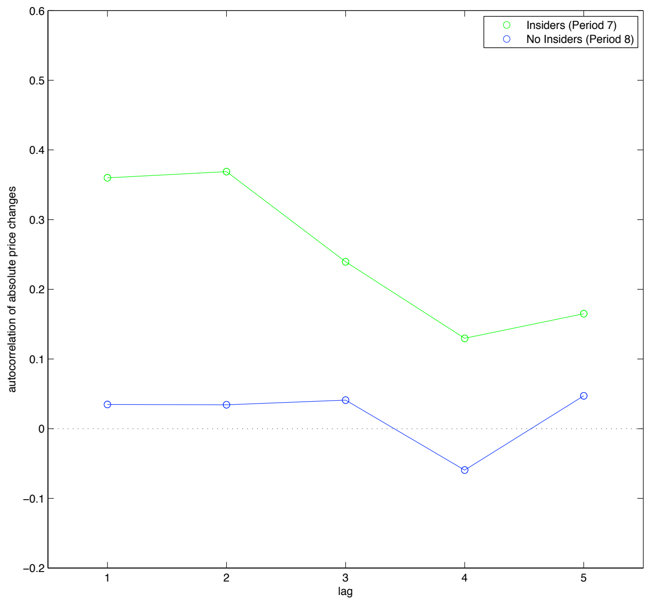

b

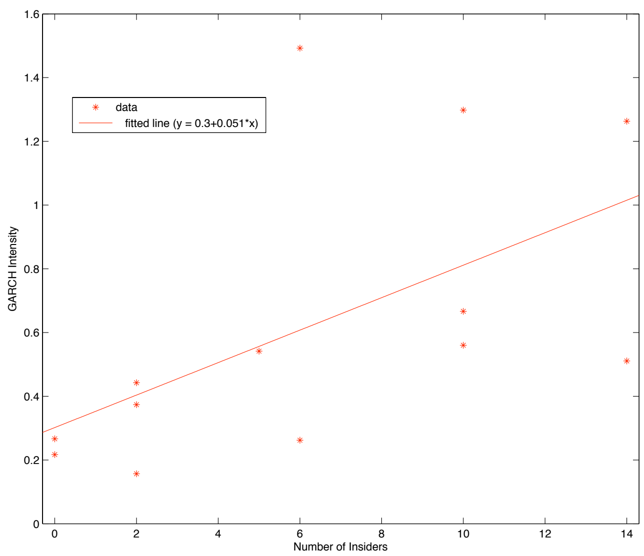

Figure 8. (a) Autocorrelation coefficients (lags 1 to 5) of absolute transaction price changes over 2s intervals in the Markets Experiment (see Figure 2), Periods 7 (insiders) and 8 (no insiders). Autocorrelation is more sizeable in Period 7. (b) Sum of absolute values of first five autocorrelation coefficients of absolute price changes ("GARCH intensity") for all periods in Markets Experiments, arranged by number of insiders; Fitted line is significant at p=0.05; R^2=0.32.

Available at: http://www.bruguier.com/pub/stockvideo.html

In case of difficulty playing the movie, we suggest using a multi-platform (Windows, Linux, MacOS X) program: http://www.videolan.org/vlc/

Video 1. Display of the trading activity. Each circle represents an offer to buy (bid, blue circle) or to sell (ask, red circle) at a certain price indicated by the number inside the circle. The diameter of the circle indicates the number of units of the stock offered. This number is the aggregate of all the offers at this price. We ordered the circles by increasing value, along one diagonal, chosen at random. The other diagonal does not carry any information and we use it to space out the circles. The circles move, grow, and shrink with the incoming orders. Every time a trade occurs, the corresponding circle (bid or ask quote that is involved in the transaction) turns green for 500ms, shrinks by the number of stocks traded, and then returns to its original color (unless no more stocks remain to be traded, in which case it disappears).

x	y	z	cluster size	t_{17}	Area
-30	-7	11	5	4.476	left insula
-14	23	39	5	4.688	frontal part of the anterior cingulate cortex
-9	41	36	22	5.380	paracingulate cortex
-9	32	45	6	4.290	frontal part of anterior cingulate cortex
17	36	43	6	6.322	frontal part of the anterior cingulate cortex
21	-10	-12	5	5.160	right amygdala

Table I. Areas with significant difference in slope coefficients to parametric regressors (insiders vs. no-insider). Standard coordinates (Talairach x,y,z) are used. We report regions with 5 or more voxels of $27mm^3$ each activated at $p<0.001$ (uncorrected) for a random effect GLM. The parametric regressor is the absolute difference between the last traded price and the 25¢. The cluster size is specified in number of contiguous voxels.

x	Y	z	cluster size	t_{17}	Area
-13	-58	-30	9	4.485	Cerebellum
-9	-65	-6	25	4.440	lingual gyrus

Table II. Areas with significant difference in slope coefficients for block regressors (insiders vs. no-insider). Standard coordinates (Talairach x,y,z) are used. Random effects, thresholded at p<0.001 (uncorrected).

Question	Answer
Consider a game played with a deck of three cards: spades, clubs, and hearts. Your goal is to identify the hearts. The cards are shuffled and displayed in a row, face down. You make your choice. The dealer then turns over one of the two remaining cards, provided it is not hearts. He then offers you the possibility to change your choice and switch to the other card that is left face down. What is the best strategy? Should you switch, stay, or does it not matter? Answer below "switch", "stay" or "either".	switch
Consider a deck of four cards: spades, clubs, hearts, and diamonds. The cards are shuffled and displayed in a row, face down. You choose one card at random and it is discarded. Then the dealer turns over two cards, chosen at random, but provided they are not hearts. Now there is only one card left unturned. If the two cards the dealer turns over are diamonds and clubs, is the probability that the remaining one is hearts more than, less than, or equal to 0.5? Answer below "more", "less" or "same".	Less
There are 8 marbles that weigh the same, and 1 marble that is heavier. The marbles are all uniform in size, appearance, and shape. You have a balance with 2 trays. You are asked to identify the heavier marble in at most 2 (two) weightings. How many marbles do you initially have to place on each tray? Input a number below.	3
Divide 100 by 1/2. Is the result more, less than or equal to 100? Answer below "more", "less", or "same".	More
Jenn has half the Beanie Babies that Mollie has. Allison has 3 times as many as Jenn. Together they have 72. Does Mollie have more than, less than, or equal to, 20 Beanie Babies? Answer below "more", "less" or "same".	More
Johnny's mother had three children. The first child was named April. The second child was named May. What was the third child's name? Type the name below.	Johnny
The police rounded up Jim, Bud and Sam yesterday, because one of them was suspected of having robbed the local bank. The three suspects made the following statements under intensive questioning. Jim: I'm innocent. Bud: I'm innocent. Sam: Bud is the guilty one. If only one of these statements turns out to be true, who robbed the bank? Type the name of the robber below.	Jim

Table III. The Mathematical (M) test. We presented subjects with seven questions in a random order. Subjects had 30 seconds to type the answer. We ignored account typos.

References

Acuna, B., J. Eliassen, J. Donoghue, and J. Sanes, 2002, Frontal and parietal lobe activation during transitive inference in humans, *Cerebral Cortex* 12, 1312-1321.

Admati, A. R., 1985, A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets, *Econometrica* 53, 629-657.

Barner, M., F. Feri, and C. R. Plott, 2005, On the microstructure of price determination and information aggregation with sequential and asymmetric information arrival in an experimental asset market, *Annals of Finance* 1, 73.

Baron-Cohen, S., T. Jolliffe, C. Mortimore, and M. Robertson, 1997, Another advanced test of theory of mind: evidence from very high functioning adults with autism or Asperger Syndrome, *Journal of Child Psychology and Psychiatry* 38, 813-822.

Baron-Cohen, S., H. A. Ring, S. Wheelwright, and E. T. Bullmore, 1999, Social intelligence in the normal and autistic brain: an fMRI study, *The European journal of neuroscience* 11, 1891-1898.

Bechara, A., and A. R. Damasio, 2005, The somatic marker hypothesis: A neural theory of economic decision, *Games and Economic Behavior* 52, 336-372.

Bhatt, M., and C. Camerer, 2005, Self-referential Thinking and Equilibrium as States of Mind in Games: fMRI Evidence, *Games and Economic Behavior* 52, 424-459.

Biais, Bruno, Peter Bossaerts, and Chester S. Spatt, 2008, Equilibrium Asset Pricing Under Heterogeneous Information, *Review of Financial Studies* Forthcoming.

Bossaerts, P., C. Plott, and W. Zame, 2007, Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiment, *Econometrica* 75, 993-1038.

Bossaerts, Peter, Charles Plott, and William Zame, 2007, Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiments, *Econometrica* 75, 993-1038.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong, 2002, Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games, *Journal of Economic Theory* 104, 137-188.

Castelli, F., F. Happe, U. Frith, and C. Frith, 2000, Movement and mind: a functional imaging study of perception and interpretation of complex intentional movement patterns, *NeuroImage* 12, 314-325.

Coricelli, G., and R. Nagel, 2008, Beauty Contest in the Brain; a Neural Basis of Strategic Thinking, *Under review*.

Crack, T. F., 2004. *Heard on the Street: Quantitative Questions from Wall Street Job Interviews* (Timothy Crack, Dunedin, New Zealand).

Critchley, H. D., C. J. Mathias, and R. J. Dolan, 2001, Neural activity in the human brain relating to uncertainty and arousal during anticipation, *Neuron* 29, 537.

Dehaene, S., E. Spelke, P. Pinel, R. Stanescu, and S. Tsivkin, 1999, Sources of Mathematical Thinking: Behavioral and Brain-Imaging Evidence, *Science* 284, 970-974.

Easley, D., and M. O'Hara, 1992, Time and the Process of Security Price Adjustment, *The Journal of finance* 47, 577-605.

Fama, Eugene F., 1991, Efficient Capital Markets: II, *The Journal of Finance* 46, 1575-1617.
Fenton-O'Creevy, Mark, Nigel Nicholson, Emma Soane, and Paul Willman, 2005. *Traders: Risks, Decisions, and Management in Financial Markets* (Oxford University Press).
Fink, G. R., P. W. Halligan, J. C. Marshall, C. D. Frith, R. S. Frackowiak, and R. J. Dolan, 1996, Where in the brain does visual attention select the forest and the trees?, *Nature* 382, 626.
Fink, G. R., P. W. Halligan, J. C. Marshall, C. D. Frith, R. S. Frackowiak, and R. J. Dolan, 1997, Neural mechanisms involved in the processing of global and local aspects of hierarchically organized visual stimuli, *Brain* 120, 1779-1791.
Gallagher, H. L., and C. D. Frith, 2003, Functional imaging of 'theory of mind', *Trends in Cognitive Sciences* 7, 77-83.
Gallagher, H. L., A. I. Jack, A. Roepstorff, and C. D. Frith, 2002, Imaging the intentional stance in a competitive game, *NeuroImage* 16, 814-821.
Gergely, G., Z. Nadasdy, G. Csibra, and S. Biro, 1995, Taking the intentional stance at 12 months of age, *Cognition* 56, 165-193.
Glosten, L., and P. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of financial economics* 14, 71-100.
Green, Jerry, 1973, Information, Efficiency and Equilibrium, Harvard Institute of Economic Research Discussion Paper.
Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the Impossibility of Informationally Efficient Markets, *The American Economic Review* 70, 393-408.
Hampton, Alan N., Peter Bossaerts, and John P. O'Doherty, 2008, Neural correlates of mentalizing-related computations during strategic interactions in humans, *Proceedings of the National Academy of Sciences* 105, 6741-6746.
Heider, F., and M. Simmel, 1944, An experimental study of apparent behavior, *American Journal of Psychology* 57, 243-249.
Heider, F., and M. Simmel, 1944, An experimental study of apparent behaviour, *American Journal of Psychology* 57, 243-259.
Lo, A. W., and D. V. Repin, 2002, The Psychophysiology of Real-Time Financial Risk Processing, *Journal of Cognitive Neuroscience* 14, 323-339.
Madhavan, Ananth, 1992, Trading Mechanisms in Securities Markets, *The Journal of Finance* 47, 607-641.
McCabe, K., D. Houser, L. Ryan, V. Smith, and T. Trouard, 2001, A functional imaging study of cooperation in two-person reciprocal exchange, *Proceedings of the National Academy of Sciences* 98, 11832-11835.
Morris, J. S., A. Ohman, and R. J. Dolan, 1998, Conscious and unconscious emotional learning in the human amygdala, *Nature* 393, 467-470.
Newman, S., P. Carpenter, S. Varma, and M. Just, 2003, Frontal and parietal participation in problem solving in the Tower of London: fMRI and computational modeling of planning and high-level perception, *Neuropsychologia* 41, 1668-1682.
Oberlechner, T., 2004, Understanding the Foreign Exchange Market through Metaphors, *British Journal of Social Psychology* 43, 133-156.

Parsons, L., and D. Osherson, 2001, New evidence for distinct right and left brain systems for deductive versus probabilistic reasoning, *Cerebral Cortex* 11, 954-965.
Phan, K. L., T. D. Wager, S. F. Taylor, and I. Liberzon, 2002, Functional Neuroanatomy of Emotion: a Meta-Analysis of Emotion Activation Studies in PET and fMRI, *NeuroImage* 16, 331-348.
Phillips, M. L., A. W. Young, S. K. Scott, A. J. Calder, C. Andrew, V. Giampietro, S. C. Williams, E. T. Bullmore, M. Brammer, and J. A. Gray, 1999, Neural responses to facial and vocal expressions of fear and disgust, *Proceedings of Royal Society of London.* 265, 1809-1817.
Plott, C. R., and S. Sunder, 1988, Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets, *Econometrica* 56, 1085-1118.
Plott, Charles R., and L. Smith Vernon, 1978, An Experimental Examination of Two Exchange Institutions, *The Review of economic studies* 45, 133-153.
Radner, Roy, 1979, Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices, *Econometrica* 47, 655-678.
Servos, P., R. Osu, A. Santi, and M. Kawato, 2002, The neural substrates of biological motion perception: an fMRI study, *Cerebral Cortex* 12, 772-782.
Singer, T., B. Seymour, J. O'Doherty, H. Kaube, R. J. Dolan, and C. D. Frith, 2004, Empathy for Pain Involves the Affective but not Sensory Components of Pain, *Science* 303, 1157-1162.
Stahl, Dale O., 2000, Rule Learning in Symmetric Normal-Form Games: Theory and Evidence, *Games and Economic Behavior* 32, 105-138.
Stone, V. E., S. Baron-Cohen, and R. T. Knight, 1998, Frontal lobe contributions to theory of mind., *Journal of Cognitive Neuroscience* 10, 640-656.
Uller, Claudia, and Shaun Nichols, 2000, Goal attribution in chimpanzees, *Cognition* 76, B27-B34.
Yoshida, W., R. J. Dolan, and K.J. Friston, 2008, Game Theory of mind, *Under review*.
Zhuanxin, Ding, W. J. Granger Clive, and F. Engle Robert, 2001, A long memory property of stock market returns and a new model, in *Essays in econometrics: collected papers of Clive W. J. Granger* (Harvard University Press).

[*] Acknowledgements: This work was supported by the NSF through grant SES-0527491 to Caltech and by the Swiss Finance Institute. Comments from Ralph Adolphs, Colin Camerer, Fulvia Castelli, John Dickhaut, John Ledyard, Charles Plott, the Editor, an Associate Editor and the referees, and seminar participants at BI (Oslo), EPFL, HEC (Paris), Northwestern University (Kellogg), the National University of Singapore, Okinawa Institute of Science and Technology, Rice University, the University of Zurich, University of Louvain, University of Texas-Austin and Dallas, the 2007 Caltech-Tamagawa Neuroscience Symposium, the 2007 Japan Neuroscience Society meetings, the 2008 meetings of the Swiss Finance Institute, the 2008 Western Finance Association meetings (especially the discussant, Heather Tookes), the 2008 International Meetings of the Economic Science Association, are gratefully acknowledged.

[1] One of the ways that we know that Theory of Mind is rather uniquely human (or shared only with higher nonhuman primates) is that tasks such as recognizing intent in an opponent's eyes engages brain structures that are evolutionary late developments (and present only in a few other primates), such as paracingulate cortex, the most frontal and medial part of the cortex.

[2] As such, ToM actually involves an aspect of learning to play games that has barely come to the attention of game theorists [exceptions include Stahl, Dale O., 2000, Rule Learning in Symmetric Normal-Form Games: Theory and Evidence, *Games and Economic Behavior* 32, 105-138., Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong, 2002, Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games, *Journal of Economic Theory* 104, 137-188.], namely, strategic sophistication. Strategic sophistication is the level of cognitive hierarchy that a player reaches. E.g., second-level thinking involves optimally responding to my opponent's reaction to his/her conjecture as to how I would play given the history of play. Building on Hampton, Alan N., Peter Bossaerts, and John P. O'Doherty, 2008, Neural correlates of mentalizing-related computations during strategic interactions in humans, *Proceedings of the National Academy of Sciences* 105, 6741-6746., Yoshida, W., R. J. Dolan, and K.J. Friston, 2008, Game Theory of mind, *Under review*. have recently suggested that ToM is not just about improving the prediction of how one's opponent changes strategy as a result of one's own actions, but actually is meant to estimate the opponent's degree of strategic sophistication. Also consistent with the idea that ToM in game play concerns strategic sophistication, activation in paracingulate cortex in the beauty contest has recently been shown to correlate with strategic sophistication [Coricelli, G., and R. Nagel, 2008, Beauty Contest in the Brain; a Neural Basis of Strategic Thinking, *Under review*.].

[3] We focused on autocorrelations of absolute values of price changes because, in field markets, persistence is known to be higher for absolute values instead of the more widely investigated squared price changes. See Zhuanxin, Ding, W. J. Granger Clive, and F. Engle Robert, 2001, A long memory property of stock

market returns and a new model, in *Essays in econometrics: collected papers of Clive W. J. Granger* (Harvard University Press).