

Exploring the Nature of “Trader Intuition”

ANTOINE J. BRUGUIER, STEVEN R. QUARTZ, and PETER BOSSAERTS*

ABSTRACT

Experimental evidence has consistently confirmed the ability of uninformed traders, even novices, to infer information from the trading process. After contrasting brain activation in subjects watching markets with and without insiders, we hypothesize that Theory of Mind (ToM) helps explain this pattern, where ToM refers to the human capacity to discern malicious or benevolent intent. We find that skill in predicting price changes in markets with insiders correlates with scores on two ToM tests. We document GARCH-like persistence in transaction price changes that may help investors read markets when there are insiders.

THIS PAPER REPORTS RESULTS FROM EXPERIMENTS meant to explore how uninformed traders 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 (as we will do here too) that uninformed traders are quite capable of quickly inferring the signals that informed traders (insiders) have about future dividends, despite anonymity of the trading process, despite a lack of structural knowledge of the situation, and despite the absence of long histories from which they can learn the market's 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 (Fama (1991)), EMH, which states that prices fully reflect all available information. Underlying EMH are the ideas 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 process. In the extreme, prices will fully reflect all available information. However, without a better understanding of how the uninformed read

*Bruguier, Quartz, and Bossaerts are at the California Institute of Technology, and Bossaerts is also at the Ecole Polytechnique Fédérale Lausanne. 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, Tania Singer, the Editor, an Associate Editor, four referees, and seminar participants at Norwegian School of Management, Haute Ecole des Hautes Etudes Commerciales (Paris), Northwestern University (Kellogg), National University of Singapore, Okinawa Institute of Science and Technology, Rice University, 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), and the 2008 International Meetings of the Economic Science Association are gratefully acknowledged.

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 here can be expressed as an attempt to better define what is meant by “trader intuition,” and to understand why some traders are better than others. Books have been written to elucidate trading intuition (Fenton-O’Creivy et al. (2005)), and correlations with specific biological markers have been discovered (testosterone level: Coates and Herbert (2008); the relative size of the index and ring fingers, an indication of a particular genetic polymorphism: Coates, Gurnell, and Rustichini (2009)). But attempts at formalizing the phenomenon have so far failed.

In contrast, our approach is methodic. After collecting the necessary trading data from purposely controlled experimental markets with and without insiders, we run a brain imaging experiment to explore human thinking during exposure to risk from insiders. The resulting data lead us to formulate a specific hypothesis about what may be at work, namely, Theory of Mind (to be defined below). Armed with this hypothesis, we design a behavioral experiment to examine whether performance in predicting prices in markets with insiders is correlated with Theory of Mind skill (and uncorrelated with other skills, in particular, mathematical and logical reasoning).

We report the results from the markets experiment that generates the data that we use in the subsequent analysis, and from the behavioral experiment that confirms the role of Theory of Mind in markets with insiders. The brain imaging experiment and its results are discussed in the Internet Appendix.¹ While they constitute an important step toward a methodic analysis of trading intuition, in the sense that, without it, our hypothesis would have amounted to pure speculation,² the technicalities involved distract from the main purpose of this paper, which is to show that trading intuition and Theory of Mind skill are strongly related.

With hindsight, it makes intuitive sense that Theory of Mind is important in markets with insiders. Let us elaborate.

¹The Internet Appendix is located on the *Journal of Finance* website at <http://www.afajof.org/supplements.asp>.

²Our study was set up as an open-ended *reverse correlation* exercise; see Hasson Uri et al. (2004). By examining brain activation during particular episodes of market replay, and armed with knowledge about the functionality of brain regions in tasks involving financial risks, one can potentially identify what subjects were thinking. We were initially looking at signals in the striatal regions of the brain as well as the anterior insula because of known correlations with changes in assessment of expected reward and risk, respectively. See, for example, Kuhnen and Knutson (2005). Activation in ToM regions came as a surprise, as these rarely activate in financial tasks unless they involve a significant strategic component. See Hampton, Bossaerts, and O’Doherty (2008).

Theory of Mind (ToM in short) is the ability to read benevolence or malevolence in patterns in one’s surroundings. ToM is thus the capacity to read intention or goal-directness, through, among other things, mere observation of eye expression (Baron-Cohen et al. (1997)), movement of geometric objects (Heider and Simmel (1944)), the moves of an opponent in strategic play (McCabe et al. (2001), Gallagher et al. (2002), Hampton et al. (2008)), or actions that embarrass others (“faux-pas”) (Stone, Baron-Cohen, and Knight (1998)). See Gallagher and Frith (2003) for further discussion.

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. 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 (2009)). From the point of view of the uninformed trader, the winner’s curse conjures up an image of potential malevolence in the trading process. Detecting this potential malevolence, then, becomes a ToM task.

Humans are uniquely endowed with the capacity to recognize malevolence as well as benevolence in their environment. ToM is human (or shared only with higher nonhuman primates); it engages brain structures that have undergone recent evolutionary expansion and reorganization, such as the paracingulate cortex, the most frontal and medial part of the cortex. In our brain imaging experiment, distinct activation in this and other ToM-related regions helps us narrow down hypotheses about the relationship between trading intuition and ToM.

The fact that markets with insiders may be exploiting a skill that (most) humans are very good at should provide a biological foundation to the plausibility of EMH. It could also explain why experiments on information aggregation in financial markets ever since Plott and Sunder (1988) have been relatively successful. The success of information aggregation experiments is in sharp contrast with, for example, simple experiments on multiperiod asset pricing, such as the infamous bubble experiments (first studied in Smith, Suchanek, and Arlington (1988)). In theory, these two types of experiments should give rise to the same type of equilibrium—the rational expectations equilibrium REE—yet experimentally equilibrium emerges robustly only in the context of information aggregation.

ToM relies on pattern recognition, something that recently has been confirmed formally, but so far only for play in strategic games (Hampton et al. (2008)). That is, humans detect malevolence or benevolence by online tracking of changes in their environment (rather than, say, logical deduction about the situation at hand). For markets, however, it is not even known whether there are patterns in the order and trade flow that would allow one to simply identify the *presence* of insiders (let alone their intentions).

Proof that such patterns exist is an important foundation for the proposition that ToM may underlie trading intuition. Therefore, we examine whether we

ourselves can find features that distinguish insider trading from normal trading in our own markets experiment. We discover that GARCH-like features (Engle (1982)) emerge when there are insiders. Specifically, autocorrelation coefficients of absolute transaction price changes in calendar time are significantly more sizeable in the presence of insiders.

The analysis in this paper is limited to thinking about and prediction in markets when there are insiders. These constitute only two of the three steps toward successful investment. We leave for future research the third step, namely, conversion of analysis and prediction into successful positions, that is, the trading itself. Indeed, superior forecasting performance does not necessarily translate into superior investments. Still, if the latter is lacking, the trading can be delegated to others who are better at it.

The remainder of this paper is organized as follows. Section I describes the markets and behavioral experiments, while briefly discussing the imaging results that identify ToM as a viable hypothesis. Section II presents the results of the markets experiment, followed by the results of the behavioral experiment. In Section III, we turn back to the markets experiment and attempt to identify whether there are patterns in the trade flow that would allow one to recognize that there are insiders, and hence on which ToM thinking could build. Section IV finishes with concluding remarks.

I. Description of the Experiments

Here, we provide descriptions of the experiments. We first run a markets experiment, for the purpose of generating order and trade flow in a controlled setting. Next, we run a brain imaging experiment, to discern how subjects judge the data, by localizing areas of the brain that are active during re-play of the markets. This leads us to identify ToM as a viable theory about the nature of trading intuition. With the ToM hypothesis in hand, we subsequently organized a behavioral experiment, where we test for correlation between subjects' ability to predict transaction prices and their generic ToM skills (we also test for correlation with mathematics and logic skills, for comparison).

We do not discuss the imaging experiment in much detail, as it provides the foundation for our hypothesis about ToM. The imaging data do constitute supportive neurobiological evidence of the behavioral findings, and without them the ToM hypothesis would have been mere speculation, but the technicalities involved in proper description of the imaging results would distract from the main point of the paper. The interested reader will find a full discussion of the imaging experiment in the Internet Appendix.

A. The Markets Experiment

Twenty subjects (undergraduate and graduate students at Caltech) participated in the markets experiment. The following situation was replicated several times, each replication being referred to as a session.

In each session, the subjects were initially endowed with notes (a risk-free asset), cash, and two risky securities, all of which expired at the end of the session. The two risky securities (“stocks”) paid complementary dividends between 0 and 50 cents: if the first security, stock X, paid x cents, then the second security, stock Z, would pay $50-x$ cents. The notes always paid 50 cents. 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, based on theory as well as prior experiments with a similar number of subjects (Bosschaerts, Plott, and Zame (2007)), prices were expected to converge to levels that equal expected payoffs; that is, risk-neutral pricing was expected to arise.

Subjects could trade their holdings for cash in an anonymous, electronic continuous open-book exchange system called jMarkets (see <http://jmarkets.ssel.caltech.edu/>). Subjects were not allowed to trade security Z, however. Consequently, risk- or ambiguity-averse subjects who held more X than Z would want to sell X to obtain a diversified (or perhaps a completely balanced) position and those who held more Z than X would need to buy X. Because there was an equal number of shares of Z and X, price pressures from trading X because of risk or ambiguity aversion were expected to cancel out.

After markets closed, liquidating dividends were paid, which, together with the remaining cash, were credited to an account. This account cumulated the earnings from each session, and at the end of the experiment subjects took home the balance in their accounts, in addition to a participation reward of \$5. All accounting was done in U.S. dollars and subjects made \$55 on average, with a range stretching from \$5 (minimum) to over \$100 (maximum).

Our trading system, jMarkets, ensured that subjects could only submit orders that, if executed, would not cause the subjects to default on their obligations at the end of a session. In calculating whether a proposed order would violate this bankruptcy constraint, the system took into account the information a subject possessed. If, for instance, a subject knew for sure that the dividend on X was not going to be above 15 cents, then we allowed the subject to take an unlimited short sale position as long as the price was above 15 cents.

In total, 13 sessions were run. Each new session started with a fresh allocation of securities and cash. As such, sessions were independent replications of the *same* situation, with the exception of one detail, namely, the information provided to subjects about the final dividend.

Specifically, in sessions that we refer 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 10 cents of the actual dividend. All subjects were always informed as to whether there were insiders. In some sessions, only the insiders knew how many insiders there were. Notice that all insiders were given the same signal. Since there were at least two insiders, insiders knew they were competing. We refer to sessions without inside information as control sessions.

Full parametrization of the markets experiment (initial endowments, signals, outcomes, etc.) can be found in the Internet Appendix. The reader can also consult the web pages through which the experiment was run:

<http://clef.caltech.edu/exp/info/>. These pages include instructions and a chronology of the public messages sent to the subjects. A copy of the instruction pages is included in the Internet Appendix. The results of the experiment are typical; the same setup has since been replicated more than 20 times, with little qualitative change in the order flow and price evolution. The reader interested in these replications is referred to Bossaerts, Frydman, and Ledyard (2009).

B. The Brain Imaging Experiment

Eighteen new subjects (undergraduate and graduate students at Caltech) were shown a replay of the 13 sessions from the markets experiment (see Section I.A), in random order, while their brains were being scanned with functional magnetic resonance imaging (fMRI). It is important to note that these subjects did *not* participate in the markets experiment. Since the fMRI experiment was held approximately 1 year after the markets experiment, we expected little contamination (in a small community like Caltech, potential subjects may talk to each other about experiments that they participate in).

The new subjects played the role of uninformed investors: while they were given the instructions of the markets experiment, they were not given any signals. At the beginning of a session, subjects were first told whether there were insiders (but never how many). Subjects then had to decide whether they would take a position of 10 units in stock X or stock Z. This feature was designed to add an element of “double blind” control: the experimenter may have known that stock X would not do well in the upcoming session, but the subject could choose stock Z instead; similarly, the subject was in control of her choice, but did not know the outcome. Subsequently, the order flow and transaction history of stock X was replayed in a visually intuitive way. During replay, subjects were only asked to push a button each time they saw a trade. As such, they could *not* change their position (say, from a position in X to a position in the complementary security, Z). At the end of the session, the liquidating dividend of the stock they had chosen was shown, and the subject was paid accordingly.

The fMRI task was deliberately kept simple. In particular, we refrained from allowing the subjects to change their trading position during replay of the order and trade flow, as we were interested in detecting, through specific patterns in brain activity, what subjects were *thinking* as the price of stock X, and hence the value of their position, went up or down. In sessions without insiders, these price movements should have been considered without consequence, as the subject was compensated based only on the final dividend on the initially chosen position. When insiders were present, however, price changes in stock X would give the subject estimates of changes in the expected final value of the chosen position.

We wanted subjects to think only about the implications of order flow and market prices for the final value of their position. If we had allowed our subjects to change positions during replay, they would have also thought about the effect of their transactions on their cash position, and in turn the value of their final

position. Since we wanted to focus our experiment on inferring value estimates from order flow and market prices, we disallowed position changes during the session.

Subjects paid a small penalty every time they missed a trade during replay of a session. This mechanism helped ensure that our subjects paid attention. Still, the bulk of total earnings came from the dividends on the positions that subjects chose before a session started. In total, the typical fMRI experiment lasted about 1 1/2 hours.

The fMRI experiment revealed increased activation in specific brain regions when insiders were present relative to when insiders were absent, and this activation differential increased as the price moved away from the unconditional expected payoff (25 cents). Such price movements should have indicated that insiders had received a signal that the value of stock X was substantially different from 25 cents. Brain regions that activated could immediately be recognized as those involved in ToM thinking (Gallagher and Frith (2003) list the regions). The most important of these is the medial paracingulate cortex, a region in the middle of the forebrain, high above the eyes.

Figure 1 displays a medial cross section of a typical human brain, from front (left-hand side) to back, onto which the significant differential activation in the paracingulate cortex is mapped (small squares; changes from dark to light correlate with increases in significance level). The localization is based on a random effects analysis of the 18 subjects' fMRI signals. Details can be found in the Internet Appendix.

Surprisingly (although this finding is corroborated in the behavioral experiment below), brain regions known to be involved in formal mathematical and logical thinking were no more activated when insiders were present than when they were not. Thus, thinking about markets when there are insiders appears to be an unequivocal ToM occupation.

These brain activation patterns prompted us to formulate the hypothesis that ToM is engaged when insiders are present. However, brain areas are generally engaged in multiple activities. As such, activation does not necessarily imply use of one specific ability. We thus turned to a behavioral experiment that we predicted would show significant correlation between, on the one hand, performance in predicting price changes when insiders are present, and, on the other hand, ToM skills as traditionally measured. We also wanted to verify the absence of significant correlation with other skills (specifically, mathematical and logical thinking), since the imaging experiment only revealed engagement of ToM regions. We now describe this experiment.

C. The Behavioral Experiment

Forty-three new subjects (undergraduate and graduate students at Caltech, Pasadena City College, and UCLA) were given a series of four tasks that were administered in random order. These subjects had participated neither in the markets experiment nor in the fMRI experiment. The four tasks were as follows.




Figure 1. Main result from the imaging experiment. Subjects' brain activation was contrasted during replay of markets with insiders against replay of markets without insiders. Shown is a cross section of the typical human brain from front (left-hand side) to back ("sagittal cross section") along the brain midline. Voxels (cubic sections of 3 mm^3) are mapped where the fMRI signal increased more intensely as a function of the difference between the transaction price and the unconditional expected payoff on stock X (25 cents) when there were insiders relative to when there were none. The significance level (p) increases from red to orange; $p < 0.001$ always; only clusters of at least five significant voxels are shown. The two clusters of significant voxels are in the paracingulate cortex. The cluster containing the voxel with the highest significance level is circled; see the Internet Appendix for precise coordinates of this voxel.

The first was a financial market prediction task ("FMP Test"), in which the order and trade flow from sessions with insiders in our markets experiment was replayed at original speed and paused every 5 seconds. During half of the pauses, we asked subjects to predict whether the last trade in the next 5 seconds was going to occur at a higher, lower, or identical price as the last trade before the pause. (When no trade occurred in the 5-second interval, the price was considered to have remained the same.) For the other half of the pauses, we reminded subjects of their predictions and informed them of their success (whether their bet had been right or not). A penalty was imposed for absence of response within a short time interval. Order and trade flow was replayed using an intuitive graphical display, discussed in more detail below.

The second task was a ToM task based on the Heider movie ("Heider Test"), a display of geometric shapes whose movements imitated social interaction (Heider and Simmel (1944)). As with the FMP task, we paused the movie every 5 seconds. For half the pauses we asked subjects to predict whether two of the shapes would get closer or not. For the other half of the pauses, we reported the outcome (whether the shapes had moved closer) and subjects' success or failure

in predicting the outcome. We ran this ToM test in the spirit of our market prediction test, namely, as a forecasting exercise where correct forecasts are rewarded and wrong forecasts are not. This is unlike the way the test is usually applied in psychology. Psychologists ask for a description of the situation and rely on verbal evidence of anthropomorphization to determine the extent to which a subject engages in ToM during replay of the Heider movies. Our variant of the test is more direct (anthropomorphization is sufficient for ToM, but not necessary), objective (verbal accounts may be misleading), and consistent with the standards of experimental economics (we paid for performance).

The third task was a ToM task based on eye gaze (“Eye Gaze Test”; Baron-Cohen et al. (1997)). A number of photographs of eye gazes were shown consecutively, and the subject is asked to pick among four possibilities that best described the mental state of the person whose eyes were shown. To facilitate this task, the subject was first given a list of words that described mental states (such as: anxious, thoughtful, skeptical, suspicious), along with a short explanation. This ToM task provided a standard test of ToM skills. In contrast with practice in psychology, we paid for performance, rewarding the subject for correct answers.

Based on our hypothesis, our conjecture was that performance in the FMP task would be correlated with scores on both the Eye Gaze and Heider tasks.

The fourth task, a mathematics task (“Mathematics Test”), consisted of a number of standard mathematics and logic questions of the type frequently used in Wall Street job interviews—see Crack (2004). For instance, one question was a variation on the Monty Hall problem, a test of understanding of Bayesian inference, which has been used elsewhere in experimental finance (Kluger and Wyatt (2004), Asparouhova et al. (2009)). A table with the seven questions used in the Mathematics Test is available in the Internet Appendix.

We added the Mathematics Test as a control: we were interested in determining whether performance in the market prediction test was correlated with ToM skills and not with other skills that arguably may also play a role in reading prices in markets with insiders.

In the financial market prediction task (FMP), we used an intuitive graphical replay of the order and trade flow. We put all the (limit) orders on the diagonal or counter-diagonal, in the form of circles. Blue circles, below the midpoint, indicated offers to buy (bids), while red circles, above the midpoint, indicated offers to sell (asks). The circles were ordered by price level. The price level itself was written inside the circle. The diameter of the circles increased with the number of units bid or asked at the corresponding price level. Whenever a trade occurred, the best bid (if a sale) or best ask (if a purchase) briefly (0.5 seconds) changed color, to green, after which the circle either shrank (if units remained available after the trade) or disappeared (if all units were traded). The circles constantly re-arranged to ensure that the best bid and ask straddled the midpoint of the screen in a symmetric way. Time remaining in the session was indicated in the unused top quadrant.

Figure 2 provides a snapshot of the graphical display (where we transform the colors into gray scale as noted in the caption). We used this display instead




Figure 2. Snapshot of the graphical replay of the order and trade flow in the financial markets prediction (FMP) task. Red circles are asks; blue circles are bids. The size of the circles increases with the number of units available. The best ask (bid) temporarily turns green when a purchase (sale) occurs. The time remaining is indicated in the top left corner (minutes : seconds : hundreds).

of the original trading interface through which subjects traded in the markets experiment because it revealed all the information without having to navigate the page, though it missed the functionality to submit orders.

II. Results

In this section, we first describe the trading data that emerged from the markets experiment, and on which the behavioral (as well as imaging) experiment was based. We subsequently discuss the results from the four performance tasks of the behavioral experiment.

A. Trading Data from the Markets Experiment

Figure 3 displays the evolution of transaction prices throughout the markets experiment. The horizontal axis denotes time; the vertical axis price level (of stock X). Vertical lines delineate sessions. Dark horizontal line segments denote final dividend level (of stock X) while light line segments denote the signal (if there were insiders). Number of insiders (I) is displayed for each period. “# K ” indicates whether everyone (“All”) or only the insiders (“Ins”) knew how many insiders there were.




Figure 3. Evolution of transaction prices (of stock X) in the markets experiment. The 13 sessions are delineated by the vertical lines; trading in a session lasted 5 minutes. Red horizontal line segments denote final liquidating dividends of stock X. Green line segments denote insider signals. Asterisks denote trade prices. In text boxes: “I” denotes the number of insiders during the session; K# denotes whether everyone (“All”) subjects or only insiders (“Ins”) knew how many insiders there were. All subjects always knew *whether* there were insiders. Stars indicate transaction prices. They are located on a grid because the tick size in the trading interface was finite (1 U.S. cent).

Trading was brisk, independent of the type of session; on average, traders entered or cancelled an offer every 0.7 seconds and one transaction took place every 3.2 seconds. In test sessions (when insiders were present), prices tended to move toward the signal, although revelation was not perfect. Closer inspection reveals a relationship between price quality (how far the final price is from the insider signal) and the proportion of informed subjects, consistent with the noisy rational expectations equilibrium REE (Grossman and Stiglitz (1976), Admati (1985)). This relationship is explored further in Bossaerts et al. (2009). In control sessions, prices tended to remain close to the competitive equilibrium (25 cents), but occasionally deviated substantially (e.g., in Session 8).

These results confirm the findings from many prior studies on information aggregation in financial markets, starting with Plott and Sunder (1988). The amplification of information through the order and trade flow is not perfect, however. But the amount of revelation is still surprising, especially because subjects do not have structural knowledge of the situation at hand to know how prices relate to signals, unlike in the theory (the rational expectations equilibrium, REE). For instance, they do not know that there is no aggregate risk, and

hence that in equilibrium pricing is as if the marginal investor is risk neutral. Yet they need this information to correctly infer information from prices.

It is clear that our subjects are rather good at inferring information from the order and trade flow, despite their lack of formal financial training. One can therefore conjecture that the situation exploits a skill that they are good at. ToM is one skill that humans are generally good at, and our behavioral experiment verified that, indeed, ToM is at work in markets with insiders. The imaging experiment suggests further that subjects did exhibit ToM (mentalized) when watching the replay of a market in which they held a stake.

B. Performance Across Tasks in the Behavioral Experiment

In the financial market prediction task, subjects were quite successful at forecasting the direction of price changes in the presence of insiders. Their forecasts were correct in approximately two-thirds of the cases on average. For comparison, randomly switching between forecasting a price increase, a price decrease, and a level price would have produced a correct forecast only 33% of the time, and a better naïve strategy, to always predict the previous outcome, would have generated a correct forecast 56% of the time. As such, subjects somehow managed to read enough information from the order flow to beat naïve forecasting rules.

There was significant variation in performance across subjects. The worst subject forecasted 46% of the cases correctly (slightly worse than the best naïve strategy), and the best subject forecasted 78% of the cases correctly. We conjectured that different levels of ToM skill explain this discrepancy in performance.

To test the above conjecture, we measure ToM skill in two ways: through the score on the Heider Test and through the score on the Eye Gaze Test. These scores provide only a rough metric of how good one is at ToM tasks, so we refrain from using them as explanatory variables in a regression of financial market prediction performance. Instead, we report correlations and their significance because correlations allow both the dependent variable (FMP task performance) and the explanatory variables (Heider or Eye Gaze test score) to be observed with error, unlike projections.

Figure 4 displays the correlation line of performance in the FMP task with scores on the Heider Test, while Figure 5 shows the same for the Eye Gaze Test (bottom panel). In both cases, the correlations are significant ($p = 0.048$ and $p = 0.023$, respectively). This confirms our hypothesis that ability to predict price changes in the presence of insiders is correlated with ToM skill.

In contrast, Figure 6 shows *no* significant correlation between performance in the financial task (forecasting price changes when there are insiders) and the score on the mathematics task (which tests mathematical and logical reasoning capacity).

Interestingly, we also do not find significant correlation either between the scores on the two ToM tests (Figure 7). It thus seems that these two tests look at different aspects of ToM, a finding that should be of interest to psychologists,




Figure 4. Relation between financial market prediction skill and ToM skill. Correlation between score on the Heider ToM Test (horizontal axis) and performance on the financial market prediction Task (vertical axis). Number of observations: 43; correlation coefficient: 0.348 ($p = 0.022$).

who generally consider the tests to be interchangeable. In our case, the lack of (significant) correlation between the scores on the two ToM tests actually is a good thing: it indirectly confirms that general intelligence or state of attentiveness cannot explain the significant correlations between scores on the FMP task and the ToM tests.

Self-reports after the behavioral experiment did not show any evidence of personalization (anthropomorphization) in the FMP task, but we found plenty of evidence pointing to personalization in the Heider Test. We did not observe significant gender differences for any test (although our subject pool was not gender-balanced: only 16 out of 43 participants were female).

The absence of significant correlation between the scores on the financial market prediction task and the mathematics test further corroborates our hypothesis that the former is a ToM task. Indeed, performance on some ToM tasks has been found to be generally uncorrelated with capacity to perform formal mathematical and logical reasoning. Specifically, through brain imaging, it has recently been found that strategic game play also constitutes a ToM task (it




Figure 5. Relation between financial market prediction skill and ToM skill. Correlation between score on the Eye Gaze ToM Test (horizontal axis) and performance on the financial market prediction task (vertical axis). Number of observations: 43; correlation coefficient: 0.303 ($p = 0.048$).

engages the usual ToM brain regions; McCabe et al. (2001), Gallagher et al. (2002)). Coricelli and Nagel (2009) show further that skill in playing the beauty contest game is not correlated with ability to do the very calculations implicit in skillful play of that game. Similarly, we find here that the ability to forecast the direction of price changes in the presence of insiders is not correlated with one's skill at formal mathematical reasoning.

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

Our finding that forecasting price changes in markets with insiders and ToM skill are related may not come as a surprise. After all, both activities concern reading the mind of an intentional source or an entity behind which there are intentional sources. In the one case, it is the market's mind that is to be read; in the other case, it is another person's (or persons') mind that is to be deciphered.

Still, formally, ToM remains a rather elusive concept. It is mostly defined only vaguely, and often in terms of specific tasks (Gallagher and Frith (2003)),




Figure 6. Relation between financial market prediction and mathematical skills. Correlation between score on the Mathematics Test (horizontal axis) and performance on the financial market prediction task (vertical axis). Number of observations: 43; correlation coefficient: 0.061 ($p = 0.699$).

or in terms of activation of particular brain regions (McCabe et al. (2001)). It is generally accepted, however, that ToM concerns pattern recognition.

In the context of strategic games, recent studies successfully identify patterns in the moves of one’s opponent on which ToM builds (Hampton et al. (2008), Yoshida, Dolan, and Friston (2008)). The import of such findings is that ToM can be concretized in terms of precise mathematical quantities that characterize an opponent’s actual play. Specifically, ToM concerns “online” or “on the fly” learning of game play. This is consistent with the proposition that ToM involves pattern recognition. Therefore, in the context of strategic games, ToM contrasts with Nash reasoning, where players would simply hypothesize that opponents choose Nash equilibrium strategies and that they would stick to them. Nash reasoning can be “offline”—it works even if one never sees an opponent’s move. Nash reasoning is also abstract—it posits only what the opponent could rationally do and how to optimally respond. Consistent with the idea that ToM and Nash thinking have little in common, brain regions that are known to be engaged in abstract mathematics do not display




Figure 7. Relation between scores on two ToM tests. Correlation between scores on the Heider ToM Test (horizontal axis) and the Eye Gaze ToM Test (vertical axis). Number of observations: 43; correlation coefficient: 0.019 ($p = 0.904$).

significant activation during game play, and, as we mention before, skill in strategic play and mathematical capabilities are uncorrelated (Coricelli and Nagel (2009)).

In our financial markets with insiders, however, it is as of yet unclear which patterns subjects could be exploiting when attempting to read the mind of the market. In fact, it has not even been established whether there are *any* patterns that distinguish markets with and without insiders. That said, our finding that subjects engage in ToM to comprehend insider trading and the view that ToM concerns pattern recognition together predict that such patterns should exist. This provides the impetus to search for them in our own markets data.

We look 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, persistence in the size of transaction price changes in calendar time provides the only statistically significant discrimination. As such, GARCH-like features appear to distinguish our sessions with and without insiders.

Specifically, we compute transaction price changes over intervals of 2 seconds.³ We follow standard practice and take the last traded price in each 2-second interval as the new price, and if there was no trade during an interval, we use the transaction price from the previous interval.⁴

We then compute the first five autocorrelations in the size (absolute value)⁵ of transaction price changes. Figure 8a plots the results for two adjacent sessions, Sessions 7 and 8. As can be inferred from Figure 2, there were 14 insiders (out of 20 subjects) in Session 7, while there were none in Session 8. The patterns in the autocorrelations of the absolute price changes in the two sessions are very different. In particular, there is substantial autocorrelation at all lags for Session 7 (when there were insiders) while there is none for Session 8 (when there were no insiders).

Figure 8b shows that this is a general phenomenon. Plotted is the sum of the absolute values of the autocorrelation coefficients for lags 1 to 5 against the number of insiders. We refer to the former as “GARCH intensity” because it measures the extent of persistence in the size of price changes. We sum the *absolute values* because closer inspection of the data reveals that autocorrelation coefficients can be significantly negative as well as positive when there are insiders. GARCH intensity increases with the number of insiders; the slope is significant at the 5% level and the R^2 , 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 thinking could build. Future research should clarify this link, which is likely to be complex. For instance, we do not find a significant correlation between GARCH intensity for a session and subjects’ performance on the financial market prediction task for the same session. Also, future research should determine whether GARCH-like features typify markets with insiders more generally, rather than just our experimental markets.

IV. Concluding Remarks

We report here how skill in forecasting price changes in markets with insiders is correlated with the general ability to detect intentionality in one’s environment, namely, ToM. This possibility was first suggested by a brain imaging experiment, whereby we contrasted activations during replay of markets with

³A trade occurred every 3.7 seconds on average, so calendar-time tick size was chosen to be slightly shorter than the average duration between trades.

⁴Note that we compute autocorrelations in calendar time. Autocorrelations are affected by stale prices because when no trade occurs, the transaction price is set equal to the last traded price. As such, autocorrelations of absolute price changes indirectly capture persistence in duration between trades as well. Still, we found no significant correlation between autocorrelation in duration between trades, on the one hand, and the number of insiders, on the other hand.

⁵We 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, Granger, and Engle (2001).





Figure 8. GARCH intensity is related to the number of insiders. (a) Autocorrelation coefficients (lags 1 to 5) of absolute transaction price changes over 2-second intervals in the markets experiment, Sessions 7 (insiders) and 8 (no insiders). Autocorrelation is more sizeable in Session 7. Vertical line segments indicate 90% confidence intervals. (b) Sum of absolute values of first five autocorrelation coefficients of absolute price changes (“GARCH intensity”) for all sessions in the markets experiment, arranged by number of insiders. The fitted line is significant at $p = 0.05$.

and without insiders: the activation in specific brain regions (and absence of activation in others) suggested that ToM may be at work. Thus, the results from the original imaging experiment and those from the behavioral experiment are consistent.

We do not find (significant) correlation between performance in the financial markets prediction task and ability to solve abstract mathematical and logical problems, nor does the imaging experiment reveal any activation in the brain regions known to be engaged in solving such problems. This finding resonates well with the extant ToM literature. Skill in playing strategic games, for instance, is uncorrelated with ability to explicitly perform the calculations that such skill implies (Coricelli and Nagel (2009)).

Our findings are of interest not only to finance. In psychology, the scope of ToM has always been confined to small-scale social interaction. Here, we demonstrate that ToM is relevant for thinking about large-scale, anonymous social structures as well. For example, ToM concerns competitive financial markets. One can envisage that it encompasses political systems as well, like voting in an election where the chance that one’s vote is pivotal is miniscule (which has left political scientists wondering why people vote at all; see, for example, Blais (2000)).

Our finding that scores on two widely accepted ToM tests are not correlated should also be of interest to the psychology community, as it suggests that ToM is not one dimensional. Indeed, our results demonstrate that forecasting prices in markets with insiders involves multiple aspects of ToM, as performance correlates with scores on *both* tests.

Note that while we set out to study “trading intuition,” a word of caution is in order. For reasons spelled out earlier, in the imaging experiment we only look at what people were *thinking* when replaying markets—subjects had to take a position before the replay, but could not change it during replay. Similarly, the behavioral experiment focuses only on *forecasting* price changes in markets with insiders—subjects were paid for the accuracy of their (directional) forecasts, but could not take positions. Of course, *thinking* about prices and *forecasting* them are integral to successful trading, but these two steps alone leave out the actual placing of orders. Trading intuition concerns not only assessment of what is going on in the market and prediction of future prices, but also submission of the right orders. Our study only considers the first two facets; future work should shed light on the third.

Our discovery that transaction prices in our sessions with insiders exhibit GARCH patterns should motivate further work. For instance, it would be interesting to know whether this is true in general. Also, we miss identification of the precise aspects of GARCH patterns that allow uninformed market participants to read the “mind of the market” (i.e., the information of the insiders), and we do not know whether the ToM brain activation that we recorded was in response to GARCH features. One could potentially get at this question by running further markets experiments like the one presented here, and recording subjects’ choices, eye gaze, and brain signals. The latter would be facilitated by the knowledge that decision neuroscience has gained in recent years about

the nature and location of brain signals related to updates of expected rewards (McClure, Berns, and Montague (2003), O'Doherty et al. (2003)), and of reward risk (Preuschoff, Quartz, and Bossaerts (2008)).

Our findings should also inspire research to improve visual representation of order and trade flow. Since humans often are better at recognizing the nature of intention in moving (animate or inanimate) objects (Heider and Simmel (1944), Castelli et al. (2000)), 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. Our (untrained!) subjects were successful in forecasting price change in the presence of insiders (on average, they performed significantly better than the best naïve strategy). One may wonder whether this success should be attributed to our using a purely graphical interface, where order and trade flows are translated into movement of circles of various sizes and colors.

Finally, the finding that markets with insiders prompt people to use a skill (ToM) that they are generally good at may explain the popularity of betting and prediction markets, where information asymmetries abound. Uninformed participants may feel confident that they will detect insider trading when it emerges. This result may also explain why people are willing to participate in markets that require advanced problem solving skills even when they know that there are others in the marketplace that have better skill (Meloso, Copic, and Bossaerts (2009)). ToM by itself cannot explain, however, why people want to participate if such markets constitute zero-sum games. Other explanations need to be invoked, such as overconfidence.

REFERENCES

- Admati, Anat R., 1985, A noisy rational expectations equilibrium for multi-asset securities markets, *Econometrica* 53, 629–657.
- Asparouhova, Elena, Peter Bossaerts, Jon Eguia, and William Zame, 2009, Cognitive biases, ambiguity aversion, and asset pricing in financial markets, Working paper, Caltech.
- Baron-Cohen, Simon, Therese Jolliffe, Catherine Mortimore, and Mary 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.
- Biais, Bruno, Peter Bossaerts, and Chester S. Spatt, 2009, Equilibrium asset pricing under heterogeneous information, *Review of Financial Studies* 22, 1503–1543.
- Blais, Andre, 2000, *To Vote or Not to Vote: The Merits and Limits of Rational Choice Theory* (University of Pittsburgh Press, Pittsburgh).
- Bossaerts, Peter, Cary Frydman, and John Ledyard, 2009, Speed of information revelation and eventual price quality in markets with insiders, Working paper, Caltech.
- Bossaerts, Peter, Charles Plott, and William Zame, 2007, Prices and portfolio choices in financial markets: Theory, econometrics, experiments, *Econometrica* 75, 993–1038.
- Castelli, Fulvia, Francesca Happe, Uta Frith, and Chris Frith, 2000, Movement and mind: A functional imaging study of perception and interpretation of complex intentional movement patterns, *NeuroImage* 12, 314–325.
- Coates, John, Mark Gurnell, and Aldo Rustichini, 2009, Second-to-fourth digit ratio predicts success among high-frequency financial traders, *Proceedings of the National Academy of Sciences* 106, 623–628.
- Coates, J. M., and J. Herbert, 2008, Endogenous steroids and financial risk taking on a London trading floor, *Proceedings of the National Academy of Sciences* 105, 6167–6172.

- Coricelli, Giorgio, and Rosemarie Nagel, 2009, Beauty contest in the brain; a neural basis of strategic thinking, *Proceedings of the National Academy of Sciences* 23, 9163–9168.
- Crack, Timothy, 2004, *Heard on the Street: Quantitative Questions from Wall Street Job Interviews* (Dunedin, New Zealand).
- Engle, Robert F., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation, *Econometrica* 50, 987–1008.
- Fama, Eugene F., 1991, Efficient capital markets: II, *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, New York, NY).
- Gallagher, Helen, and Chris Frith, 2003, Functional imaging of “theory of mind”, *Trends in Cognitive Sciences* 7, 77–83.
- Gallagher, Helen, Anthony Jack, Andreas Roepstorff, and Chris Frith, 2002, Imaging the intentional stance in a competitive game, *NeuroImage* 16, 814–821.
- Green, Jerry, 1973, Information, efficiency and equilibrium, Harvard Institute of Economic Research Discussion Paper.
- Grossman, Sanford, and Joseph Stiglitz, 1976, Information and competitive price systems, *American Economic Review* 66, 246–253.
- 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.
- Hasson, Uri, Yuval Nir, Ifat Levy, Galit Fuhrmann, and Rafael Malach, 2004, Intersubject synchronization of cortical activity during natural vision, *Science* 303, 1634–1640.
- Heider, Fritz, and Marianne Simmel, 1944, An experimental study of apparent behavior, *American Journal of Psychology* 57, 243–249.
- Kluger, Brian D., and Steve B. Wyatt, 2004, Are judgment errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem, *Journal of Finance* 59, 969–997.
- Kuhnen, Camelia, and Brian Knutson, 2005, The neural basis of financial risk taking, *Neuron* 47, 763–770.
- McCabe, Kevin, Daniel Houser, Lee Ryan, Vernon Smith, and Theodore Trouard, 2001, A functional imaging study of cooperation in two-person reciprocal exchange, *Proceedings of the National Academy of Sciences* 98, 11832–11835.
- McClure, Sam, Gregory Berns, and P. Read Montague, 2003, Temporal prediction errors in a passive learning task activate human striatum, *Neuron* 38, 339–346.
- Meloso, Debrah, Jernej Copic, and Peter Bossaerts, 2009, Promoting intellectual discovery: Patents versus markets, *Science* 323, 1335–1339.
- O’Doherty, John, Peter Dayan, Karl Friston, Hugo Critchley, and Ray Dolan, 2003, Temporal difference models and reward-related learning in the human brain, *Neuron* 38, 329–337.
- Plott, Charles, and Shyam Sunder, 1988, Rational expectations and the aggregation of diverse information in laboratory security markets, *Econometrica* 56, 1085–1118.
- Preuschoff, Kerstin, Steven R. Quartz, and Peter Bossaerts, 2008, Human insula activation reflects risk prediction errors as well as risk, *Journal of Neuroscience* 28, 2745–2752.
- Radner, Roy, 1979, Rational expectations equilibrium: Generic existence and the information revealed by prices, *Econometrica* 47, 655–678.
- Smith, Vernon, L., Gerry L. Suchanek, and W. Williams Arlington, 1988, Bubbles, crashes, and endogenous expectations in experimental spot asset markets, *Econometrica* 56, 1119–1151.
- Stone, Valerie, Simon Baron-Cohen, and Robert T. Knight, 1998, Frontal lobe contributions to theory of mind, *Journal of Cognitive Neuroscience* 10, 640–656.
- Yoshida, Wako, Ray J. Dolan, and Karl J. Friston, 2008, Game theory of mind, *PLoS Computational Biology* 4, e1000254.
- Zhuanxin, Ding, Clive W.J. Granger, and Robert F. Engle, 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, Cambridge, MA).