JENA ECONOMIC RESEARCH PAPERS

2012 – 057

Judgmental Overconfidence and Trading Activity

by

Gerlinde Fellner
Sebastian Krügel

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a joint publication of the Friedrich Schiller University and the Max Planck Institute of Economics, Jena, Germany. For editorial correspondence please contact markus.pasche@uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

Max Planck Institute of Economics
Kahlaische Str. 10
D-07745 Jena
www.econ.mpg.de

Judgmental Overconfidence and Trading Activity*

Gerlinde Fellner† and Sebastian Krügel‡

October 12, 2012

Abstract

We investigate the theoretically proposed link between judgmental overconfidence and trading activity. In addition to applying classical measures of miscalibration, we introduce a measure to capture misperception of signal reliability, which is the relevant bias in the theoretical overconfidence literature. We relate the obtained overconfidence measures to trading activity in call and continuous experimental asset markets. Our results confirm prior findings that classical miscalibration measures are not related to trading activity. However, misperception of signal reliability is significantly linked to trading volume, particularly in the continuous market. In addition, we find that men trade more than women at high levels of risk aversion, but the gender trading gap vanishes as risk aversion lessens. The reason is that the trading activity of women seems to be more sensitive to risk attitudes than that of men.

JEL classification: D03; C91; G12;

Keywords: Overconfidence, Trading activity, Signal perception

*We are indebted to Martin Beck for research assistance and want to thank the participants of the IMPRS Workshop 2010 in Ringberg and the ESA world meeting 2010 in Copenhagen as well as seminar participants at the Frankfurt School of Finance & Management, the Max Planck Institute of Economics, the University of Cologne and Utrecht University for valuable comments.

†Ulm University, Institute of Economics, gerlinde.fellner@uni-ulm.de

‡Max Planck Institute of Economics, International Max Planck Research School 'Uncertainty' and Ulm University, kruegel@econ.mpg.de

1 Introduction

Trading activity on financial markets appears to be extraordinarily high. De Bondt and Thaler (1995) claim that the observed level of trading "is perhaps the single most embarrassing fact to the standard finance paradigm." This motivated several researchers in behavioral finance to extend traditional market models by plausible psychological biases, among which overconfidence is often viewed as the most promising to explain the "trading puzzle" (see, e.g., De Bondt and Thaler, 1995). In the theoretical studies, overconfidence is always modeled as a biased belief about the precision of private information (e.g., Odean, 1998, Kyle and Wang, 1997, Benos, 1998, Caballé and Sákovics, 2003). An overconfident investor overestimates the precision of her private information, therefore overweights this information when she updates her beliefs, and, as a consequence, ends up with a biased posterior belief about the value of an asset. Ultimately, this will lead to more trade.

In the present study, we conduct a test of the overconfidence hypothesis. Before we outline how our inquiry differs from previous empirical studies, it is important to note that the recent psychological literature distinguishes between three distinct types of overconfidence: (i) judgmental overconfidence (i.e., overestimating the precision of one's judgment), (ii) self-enhancement biases (i.e., positive self-illusions such as the better-than-average effect and illusions of control), and (iii) optimism with respect to societal risks (e.g., Hilton et al., 2011).[1] The modeled bias in the behavioral finance literature clearly falls into the first category. In order to conduct a test of the theoretically proposed link, an explicit measure is needed for the judgmental type of overconfidence.

To the best of our knowledge, there are three existing studies that attempt to test the effect of overconfidence on trading directly.[2] First, there is a study by Biais et al. (2005) in

[1] For a slightly different categorization and terminology, see Moore and Healy (2008).

[2] There are several other empirical studies which provide evidence for the trading effect of overconfidence (e.g., Barber and Odean, 2001, Chuang and Lee, 2006, Statman et al., 2006). However, all of these studies do not measure overconfidence directly and therefore need to rely on proxies for overconfidence such as gender or success. Thus, these studies can only provide suggestive evidence for the trading effect of overconfidence, which only holds under certain auxiliary assumptions. Direct tests of behavioral finance models (i.e., correlating psychological measurements with the relevant economic variables) are much more meaningful and desirable (Glaser and Weber, 2007).

which judgmental overconfidence is assessed through a so-called interval production task with respect to general knowledge. Such tasks are well established in cognitive psychology (see, e.g., Lichtenstein et al., 1982, Klayman et al., 1999): individuals have to state confidence intervals for numerical answers to several knowledge questions.[3] The common result of such tasks is that individuals' confidence intervals are too narrow. This indicates that individuals overestimate the precision of their knowledge, a phenomenon that is usually called "miscalibration." In a series of experimental asset markets, Biais et al. find, however, no relation between miscalibration scores and trading activity.

Glaser and Weber (2007) confirm this finding. In their study, they combine real trading data from investors of a German online broker with individual overconfidence scores, which they obtained through an Internet survey. To elicit judgmental overconfidence, Glaser and Weber also employ the interval production task and measure miscalibration with respect to knowledge and volatility estimates in stock market forecasting. However, neither of their miscalibration measures predicts trading activity. They conclude that "[m]easures of miscalibration are, contrary to predictions of overconfidence models, unrelated to measures of trading volume."

A different result is reported in a study by Deaves et al. (2009). Like Biais et al. (2005), Deaves and coauthors first elicit miscalibration from their subjects, based on several knowledge questions. Subsequently, the subjects participate in an experimental asset market. In line with the theory, Deaves et al. find that higher levels of miscalibration indeed predict higher levels of trading volume. However, by design, their measure of miscalibration is confounded with the better-than-average effect, which clearly is a different type of overconfidence from that modeled in the finance literature. Consequently, one cannot interpret their findings as confirming evidence in favor of the overconfidence models.[4]

As it stands, the existing evidence speaks against the theoretically derived link between judgmental overconfidence and trading volume. Does this mean that models invoking judgmental overconfidence as a reason for high trading volume have to be rejected? We argue that the lack of empirical support for the overconfidence effect on trading may be rooted in a

[3]Note that the interval production method can hardly be incentivized.

[4]Glaser and Weber (2007) also make this point. See their footnote 45.

discrepancy between modeling and measuring overconfidence (see Fellner and Krügel, 2012). Whereas the theoretical overconfidence literature models the perception of signal reliability, the measurement of overconfidence in empirical studies relies on calibration scores with respect to the misperception of own knowledge and/or time series volatility. In line with some other studies (e.g., Glaser and Weber, 2007), Fellner and Krügel (2012) find that miscalibration in a knowledge task is associated with miscalibration in a time series forecasting task, suggesting that both tasks expose the same underlying judgmental bias. However, misperception of signal reliability seems to be a distinct bias.

In the present study, we build on these results and undertake a new test of the predictions of the overconfidence models. In addition to the usual measures of miscalibration used in previous empirical studies, we also capture individuals' perception of signal reliability. To this end, a prediction task is used in which subjects have to forecast the realization of a random variable based on a noisy signal over many rounds. Subjects know that the underlying distribution of the noise term is kept constant across rounds and that the a priori signal quality is therefore the same in each round. For each subject we then regress the predictions on the corresponding signals. By this procedure, we obtain an individual measure of the weighting of information which captures the perception of signal reliability.

This proposed measure of overconfidence has several advantages over the miscalibration measures used so far in empirical tests of the overconfidence hypothesis. First and foremost, it captures the judgmental bias incorporated in the overconfidence trading models most closely. Second, the underlying task can be easily incentivized and the overconfidence measure does not rely on pure survey questions. Third, the proposed measure is inferred from actual behavior and "it is quite possible that while individuals are not able to *communicate* probabilistic assessments well, they are able to *incorporate* them into their decisions" (Kogan, 2009, p.1893).

Our paper makes another important contribution to the existing literature by investigating the link between overconfidence and trading volume while additionally controlling for risk attitudes and gender effects. Men have frequently been found to be more overconfident than women, although this effect seems to be task dependent (Lundeberg et al., 1994). Still, higher overconfidence is assumed to account for higher trading activity by men (see Barber and Odean, 2001) when, in fact, it is possible that this effect is driven by gender differences in risk attitudes (Fellner and Maciejovsky, 2007). Surprisingly, previous studies on overconfidence and trading have largely neglected the possible interaction of these aspects.

To gather trading behavior, we employ a series of experimental asset markets after obtaining the relevant psychological measures. We implement a call market and a continuous double auction. In line with previous empirical studies, we expect that miscalibration scores based on a general knowledge task and a time series forecasting task are unrelated to individual trading activity. In contrast to this, we expect that our proposed measure regarding the perception of signal reliability predicts trading activity as hypothesized by the overconfidence models.

In the next section, we illustrate the design and procedure of our experiment. Section 3 outlines the hypotheses and section 4 presents the results. In the last section, we summarize our findings and conclude.

2 Experimental Design and Procedure

The experiment consisted of five stages which were conducted in a fixed order. The first two stages served the purpose of eliciting the two measures of miscalibration that have been used to capture judgmental overconfidence in previous studies. The third stage served to measure the risk attitude of our subjects. In the fourth stage, we employed the signal-based prediction task to obtain the new measure of overconfidence regarding subjects' perception of signal reliability. Finally, stage five contained the experimental asset markets. This stage was again subdivided into two different market phases.

Subjects received written instructions for the first two stages at once and for each subsequent stage only after all subjects of the same session had completed the prior stage. The experiment was fully computerized using the software z-Tree (Fischbacher, 2007), and all subjects were recruited for participation using ORSEE (Greiner, 2004). In total, 168 students from Friedrich Schiller University Jena participated in the experiment, 85 men and 83 women with an average

age of 23.4 years (SD=2.9).[5] Each session lasted for about three hours, and earnings in the experiment ranged from 13.00 € to 42.80 € with an average of 23.10 €. In the following, we describe the stages in more detail.

2.1 Stage one: General knowledge questions

In the first stage, subjects had to state confidence intervals for several knowledge questions requiring a numerical answer. This method is well established in psychology and has frequently been used in behavioral and experimental economics to elicit judgmental overconfidence (e.g., Biais et al., 2005, Glaser and Weber, 2007, Deaves et al., 2009). In our experiment, subjects were confronted with 10 almanac questions, all of which can be found in the Appendix together with the correct answers. One of the questions was, for instance:

What is the length of the River Nile in km?

To obtain a confidence interval for such a question, subjects were asked to state a lower and an upper bound to be 90% sure that the correct answer lay within the interval. Overconfidence is captured by a *calibration index* which was calculated for each subject as the number of incidents where the correct answer lies outside the stated interval. An index of 1 indicates well-calibration, 0 reflects underconfidence and values from 2 to 10 (increasing) overconfidence.[6]

2.2 Stage two: Time series forecasting

The task in the second stage was methodologically similar to the first one because it also employed the interval production method. However, in this task, subjects were asked to state 90% confidence intervals as forecasts for 10 time series. All time series were pre-generated over 24 periods according to an autoregressive, moving average process with one MA and one AR

[5]The vast majority of our subjects were students of economics and business administration. The fields of study of the remaining subjects were mainly mathematics, physics, or computer sciences.

[6]Since we asked for 90% confidence intervals, there is an obvious asymmetry in the possibility to identify over- and underconfidence. However, we used the same method as many previous studies to be able to relate our results to them (e.g., Biais et al., 2005, Glaser and Weber, 2007, Hilton et al., 2011, Klayman et al., 1999, Russo and Schoemaker, 1992).

term. For each time series we used different parameters for the MA and AR term, but all had a common starting value of 200. A trend component was not included. The so generated time series constitute an "ideal" forecasting environment and have frequently been used in forecasting research (see, e.g., Lawrence and O'Connor, 1992, 1993).

Subjects were presented the first 20 periods of each time series and asked to forecast the value in period 24 by stating an upper and lower bound to be 90% sure that the true value would fall within the interval. Figure A.2 gives an overview of all 10 time series where subjects were, of course, only shown the solid lines from period 1 to 20. The dashed horizontal lines indicate the 90% confidence interval of the realization in period 24. The instructions for this task made clear to the subjects that all time series were computer generated such that it was impossible to recognize price patterns of real assets.

As in task 1, overconfidence is captured through a *calibration index*, which was calculated for each subject as the number of incidents where the true value in period 24 lay outside the stated interval. An index of 1 indicates well-calibration, 0 reflects underconfidence and values from 2 to 10 (increasing) overconfidence.

Note that it is not easily possible to provide performance-based incentives for the production of confidence intervals in tasks 1 and 2. Similar to most other studies, each of our subjects therefore received flat monetary incentives of 3.00 € for finishing the first two stages.

2.3 Stage three: Risk attitude

The risk attitude elicitation method follows the procedure suggested by Holt and Laury (2002). In 11 repetitions, subjects had to choose between two binary lotteries X and Y. Table 1 shows the list of the 11 lottery choices where all amounts are displayed in euro. While the prizes of both lotteries, $\overline{x}$ and $\underline{x}$ as well as $\overline{y}$ and $\underline{y}$, remain constant, the probabilities of the high prizes $p(\overline{x})$, $p(\overline{y})$ increase from choice 1 to 11 (with $p(\underline{x}) = 1 - p(\overline{x})$ and $p(\underline{y}) = 1 - p(\overline{y})$). A risk-neutral individual, who decides only upon the expected value of lotteries, would choose lottery X five times and then switch to lottery Y at the sixth choice. An earlier switch from lottery X to Y indicates risk-seeking behavior, a later switch risk-averse behavior. According to the switching behavior, a *risk index* for each individual is created in the following way: An individual who switched from lottery X to lottery Y at choice 11 is assigned a risk index of 0,[7] an individual who switched to lottery Y at the tenth choice is assigned a risk index of 1 and so on. The so-generated *risk index* runs from 0 to 9, where values from 0 to 4 indicate risk aversion, 5 indicates risk neutrality, and values from 6 to 9 risk-loving behavior.[8]

Table 1: Lottery choices for elicication of risk attitudes

No.	Lottery X			Lottery Y		
	$p(\overline{x})$	$\overline{x}$	$\underline{x}$	$p(\overline{y})$	$\overline{y}$	$\underline{y}$
1	0.0	2	1.6	0.0	3.85	0.1
2	0.1	2	1.6	0.1	3.85	0.1
3	0.2	2	1.6	0.2	3.85	0.1
4	0.3	2	1.6	0.3	3.85	0.1
5	0.4	2	1.6	0.4	3.85	0.1
6	0.5	2	1.6	0.5	3.85	0.1
7	0.6	2	1.6	0.6	3.85	0.1
8	0.7	2	1.6	0.7	3.85	0.1
9	0.8	2	1.6	0.8	3.85	0.1
10	0.9	2	1.6	0.9	3.85	0.1
11	1.0	2	1.6	1.0	3.85	0.1

At the end of the experiment, one of the 11 cases was selected at random, and subjects were paid out according to their choice of alternative X or Y and the randomly determined state that occurred for the respective lottery. It is important to note that prior to the risk elicitation stage, all subjects had gained the same earnings since the first two tasks were paid based on a flat fee. Thus, we avoided any variance in previous payoffs across subjects that could possibly influence risk attitudes. For this reason we chose to conduct the risk elicitation stage prior to measuring overconfidence in signal perception since the latter task has performance-based incentives.

[7]Note that at the eleventh choice lottery Y dominates lottery X. Thus, individuals who did not switch to lottery Y at this choice could not be assigned any risk index.

[8]Subjects who chose the dominated lottery Y over lottery X at choice one could not be assigned a risk index for the same reasons as above. Rather than risk preferences, lottery choices one and eleven reveal subjects' (mis)understanding of the task.

2.4 Stage four: Signal-based predictions

The fourth stage aimed at eliciting subjects' overconfidence in signal-based predictions, and should thus reflect the kind of overconfidence that models in economics and finance exclusively focus on. The task borrows a methodology from the psychological literature on so-called single-cue probability learning, where subjects have to predict the outcome x of a random variable $\tilde{x}$ based on a noisy signal $\tilde{s}$ over many rounds. The random variable $\tilde{x}$ is normally distributed with $N(\bar{x}, \sigma_x^2)$. Signals $\tilde{s}$ are determined by the true value plus noise, formally $\tilde{s} = \tilde{x} + \tilde{e}$, where $\tilde{e}$ is a random error term that is distributed according to $N(0, \sigma_e^2)$, and $\tilde{x}$ and $\tilde{e}$ are mutually independent. After receiving signal s, a rational decision maker updates her belief about the outcome x according to

$$E[\tilde{x}|\tilde{s} = s] = \frac{\sigma_e^2 \bar{x} + \sigma_x^2 s}{\sigma_x^2 + \sigma_e^2} \tag{1}$$

which corresponds to the linear least-square predictor of $\hat{x}$, given signal s:

$$\hat{x}[s] = \bar{x} + r_{x,s} \frac{\sigma_x}{\sigma_s}(s - \bar{x}) \tag{2}$$

where

$$r_{x,s} = \frac{\sigma_x}{\sigma_s} = \frac{\sigma_x}{\sqrt{\sigma_x^2 + \sigma_e^2}} \tag{3}$$

is the correlation between outcome and signal. Overestimating signal precision, i.e., underestimating the error variance σ_e^2, will lead to a prediction slope that is too steep, which, in turn, indicates that the private signal has been overweighted. This bias is exactly what behaviorally inspired economic models generally refer to as overconfidence.[9]

In the experiment, all subjects had to perform the signal-based prediction task over 60 rounds. The random variable $\tilde{x}$ was drawn from a normal distribution with $N(585, 50^2)$, and

[9] Economic models formally introduce a trader-specific overconfidence parameter, K_i, which captures trader i's perception of the signal precision, i.e., $\tilde{s_i} = \tilde{x} + K_i \tilde{e_i}$. Trader i is overconfident if $0 \leq K_i < 1$ such that she overweights her signal and forms a posterior belief $E[\tilde{x}|\tilde{s_i} = s_i] = \frac{K_i^2 \sigma_e^2 \bar{x} + \sigma_x^2 s_i}{\sigma_x^2 + K_i^2 \sigma_e^2}$. In the extreme case where $K_i = 0$, the overconfident trader i has a posterior belief about the asset value which is simply equal to the signal s_i. Thus, in general, the posterior beliefs of overconfident traders are more dispersed which, in turn, increases trading volume (see, e.g., Odean, 1998, Kyle and Wang, 1997).

the random error term $\tilde{e}$ from a normal distribution with $N(0, 50^2)$.[10] Subjects were informed that the chosen distributional properties were constant across rounds such that the a priori signal quality was the same in each round. Before the first prediction, subjects received a list of ten random draws of $\tilde{x}$ and corresponding signals $\tilde{s}$. They were thus able to draw some inferences about the reliability of the signals prior to their first prediction. In each round, subjects then received a signal s and had to predict x, knowing that the signal is a non-perfect but unbiased indicator of value x.[11] After each prediction, the true outcome value x was revealed and individuals moved on to the next prediction round. At the end of the task, one of the 60 rounds was randomly selected and all subjects were paid according to the accuracy of their prediction in this round, based on a linear scoring rule: for every integer that a subject's prediction deviated from the true outcome value, 1.5 euro cents were subtracted from a lump sum payment of 6.00 €.[12] Payoff tables with several examples were added to the instructions to illustrate the incentive scheme.

Following different strands of literature in economics and psychology, we implemented two treatments for the signal-based prediction task which differed in the prior information about the distribution of the random variable $\tilde{x}$. In the *Info* treatment, subjects received all information about the chosen distributional properties except for the error variance σ_e^2 of the signal. To ensure an appropriate understanding of the normal distribution of the random variable $\tilde{x}$, a chart of 1,000 random realizations from the truncated normal distribution was displayed in the instructions. This treatment most closely captures the overconfidence models in the economics literature: agents are assumed to know the underlying distribution of the variable of interest, e.g., the value of the asset. The key information they misjudge is the precision of their private signal. In the *No-Info* treatment, on the other hand, subjects were not informed about the

[10] For reasons of experimental practicality, the distributions were truncated at both ends at four standard deviations. The actual values x, e.g., were thus restricted to the range of 385 to 785, which was also known to the subjects. To ensure comparability between sessions, values x and e for all 60 rounds were pre-generated and kept constant across subjects.

[11] For more details on instructions, see the Appendix.

[12] Using a quadratic scoring rule instead of the simple linear rule would be incentive compatible, at least for risk-neutral decision makers. However, the high number of risk-averse participants according to stage 3 undermines the usefulness of a quadratic scoring rule and, additionally, subjects would be easily overburdened by the payoff rules.

distributional properties of the outcome variable. They only knew the range of possible values for $\tilde{x}$. This corresponds to the standard procedure of the psychological studies on single-cue probability learning (e.g., Czaczkes and Ganzach, 1996, Ganzach, 1993, 1994), from which we borrowed the methodology for this task. Note that in the *No-Info* treatment two sources might contribute to an overweighting bias: (1) underestimating the error variance σ_e^2, just as in the *Info* treatment, and (2) overestimating the variability (i.e., variance) of the distribution of the outcome variable $\tilde{x}$.

Of the 168 subjects, 88 participated in the *Info* and 80 in the *No-Info* treatment. As a measure of overconfidence, we obtained a prediction slope for each subject by regressing the individual predictions on the corresponding signals. The normative prediction slope, which is calculated by regressing the true outcome values on the corresponding signals, was equal to 0.5.[13] A prediction slope higher than 0.5 therefore indicates overweighting of private information and, thus, overconfidence in the reliability of signals. A prediction slope lower than 0.5 indicates underweighting of private information and, thus, underconfidence in the reliability of signals. An explicit advantage of this method is, therefore, that it equally allows for over- and underconfidence.

2.5 Stage five: Experimental asset markets

In order to gain data on individual trading behavior, we conducted fully computerized experimental asset markets in stage five. Each market was formed of eight randomly assigned subjects. With 168 participants in the experiment, a total of 21 markets were implemented. All subjects received detailed written instructions about the market procedure, the traded asset, possible earnings, and the design of the trading screen.[14] All subjects had sufficient time to read the instructions, and remaining questions were answered privately.

In each market, all eight subjects could trade an asset with unknown liquidation value $\tilde{v}$

[13]The values for $\tilde{x}$ and $\tilde{e}$ were randomly generated with the constraint that the chosen distributional properties would be preserved within the first, second, and third block of 20 rounds to investigate potential learning effects (see Fellner and Krügel, 2012). This property also ensures that the normative slope based on the sample equals the theoretical normative slope based on equation (2).

[14]The full set of instructions as well as the trading screen can be found in the Appendix.

drawn from a normal distribution with $N(120, 15^2)$.[15] Before a trading round started, each market participant i received a private signal s_i, which was randomly determined according to $\tilde{s}_i = \tilde{v} + \tilde{e}_i$, where $\tilde{e}_i \sim N(0, 13^2)$.[16] All subjects knew the distribution of $\tilde{v}$,[17] how the signal was determined and that the expected value of the error term $\tilde{e}_i$ was 0. However, as assumed in the overconfidence models, they were not informed about the variance of the error term. Before the experimental asset markets started, each subject received a list of 10 random draws for the asset value $\tilde{v}$ and corresponding signals $\tilde{s}$ to gain an impression of the reliability of the signals. This list was the same for all subjects.

At the beginning of each round, subjects were endowed with 10,000 ECU (experimental currency units) and 20 shares of the asset. Borrowing and short selling was not allowed. In total, 10 independent trading rounds with re-initialized endowments and fixed groups were conducted. All asset values as well as all eight signals per round were randomly determined prior to the experimental sessions and held constant across markets. The signals of a round were generated with the constraint that the average signal conveyed the true asset value, which was also told to subjects.

Each trading round consisted of two phases: an opening call market followed by a continuous double auction.[18] In the call market, subjects could transmit limit orders by specifying a minimum selling price together with the maximum number of shares they wished to sell and/or a maximum buying price together with the maximum number of shares they wished to buy. After all market participants had submitted their orders, an aggregate supply and demand curve was constructed and the market clearing price that maximized trading volume was set. Each participant was then informed about the market clearing price, their own concluded trades at that price, and their current amount of cash and shares of the asset. An advantage of the call market phase is to collect independent information about the willingness to buy and sell

[15] Here, too, the distribution was truncated at both ends at four standard deviations so that actual values v were restricted to a range of 60 to 180. All subjects were, of course, informed about the truncation.

[16] Again, the normal distribution of the error term was truncated at both ends at four standard deviations.

[17] To ensure an appropriate understanding of the distribution of the asset value, a chart of 1,000 random realizations from the truncated normal distribution was displayed in the instructions.

[18] An opening call auction followed by continuous trading is common practice on many stock exchanges like the NYSE.

the asset before any interaction between market participants in a double auction takes place.

After the call market phase, subjects could trade on a computerized continuous double auction market with an open limit order book for a total of 120 seconds. At the beginning of each double auction, the limit order book was empty. Traders could then submit limit asks or bids or accept standing limit orders submitted by others. The size of each order was limited to one share, and full transparency was provided with respect to concluded trades. However, trading was anonymous and no identification codes were posted in the trading phase.

At the end of each trading round, subjects received information about the true asset value, the value of their initial portfolio, their concluded trades, the value of their final portfolio, and their earnings in the current round. The value of the final portfolio was calculated by summing up the remaining amount of a subject's cash as well as the number of asset shares at the end of the round multiplied by the true asset value. The value of the initial portfolio was calculated analogously using the cash and asset share endowments at the beginning of the trading round. The difference between the values of the final and initial portfolio determined earnings of one round.

At the end of stage five, one of the ten trading rounds was randomly selected for payment, and ECUs earned in that round were converted to euro at a rate of 25:1. Since trading profits could be negative, all subjects additionally received a lump sum fee of 10.00 € for participating in the experimental asset markets. Trading gains were added and trading losses subtracted.[19]

Prior to the first trading round in stage five, subjects could practice trading in the double auction market in two training rounds. All subjects were informed that the asset value in the training rounds was set to the mean value of 120 and that gains or losses were, of course, irrelevant for their earnings in the experiment. They were also encouraged to trade during those two rounds to understand the trading procedure and become acquainted with the trading screen.

[19]In the unlikely event that subjects were still in a loss position even after the lump sum fee of 10.00 € was added to their trading profits, they knew that they could compensate losses of stage five by working on additional tasks after the experiment. However, this never happened.

3 Hypotheses

According to the predictions of the overconfidence models, individual trading activity should increase with increasing judgmental overconfidence. Previous empirical studies have not been supportive of this conjecture. This lack of empirical support could be due to the fact that the previously employed measures of miscalibration do not capture the modeled overweighting bias that revolves around the misperception of signal reliability (see also Fellner and Krügel, 2012). Based on the theoretical models and previous empirical findings, we formulate our two main working hypotheses.

Hypothesis 1 *Miscalibration, i.e., overconfidence with respect to general knowledge and time series predictability, is not related to trading activity.*

Hypothesis 2 *Overconfidence in the reliability of signals is positively related to trading activity.*

Our experimental setup allows to test these two main hypotheses while controlling for possible moderating factors on trading like risk attitude or gender. Both have been found to affect trading behavior of individuals. While Fellner and Maciejovsky (2007), for instance, show that higher degrees of risk aversion are linked to lower degrees of trading activity, several other studies find that men trade more than women (e.g., Barber and Odean, 2001, Biais et al., 2005, Grinblatt and Keloharju, 2009). In addition, gender differences in risk attitude have been extensively debated in the previous literature. A fairly robust finding is that women are more risk averse than men (e.g, Weber et al., 2002, Eckel and Grossman, 2008). Putting these findings together leads to the conjecture that gender differences in trading might be due to differences in risk attitudes between men and women. We thus formulate two supplement hypotheses:

Hypothesis 3 *The higher the individual degree of risk aversion, the lower the trading volume.*

Hypothesis 4 *Gender differences in trading disappear, once we control for risk attitude.*

As for gender differences in overconfidence, results are less clear-cut. At least in domains of ambiguous or no feedback, men have been found to be more overconfident than women (e.g., Lundeberg et al., 1994). This alleged gender difference in overconfidence has been put forward as a cause for the empirical observation that men trade more than women (Barber and Odean, 2001). It is thus an open question whether differences in overconfidence between men and women account for differences in trading activity. In other words, we conjecture:

Hypothesis 5 *Men are, on average, more overconfident than women.*

Hypothesis 6 *Gender differences in trading disappear, once we control for the degree of overconfidence.*

4 Results

We first present summary statistics of all overconfidence scores, analyze possible gender differences, and report the findings on the relation between our overconfidence measures.[20] We then describe the data in terms of individual trading activity and start with a simple bivariate analysis to address the impact of overconfidence on trading as well as possible gender differences and the effect of risk attitude. In the last subsection, we report the results of several regressions to investigate the effect of overconfidence on trading activity in more detail.

4.1 Overconfidence

Panel A of Table 2 gives a descriptive overview of the overconfidence measures based on the three different tasks. In the general knowledge and time series forecasting task, an individual index of 1 indicates well-calibration, values above 1 reflect overconfidence, and values below 1 underconfidence. Panel A shows that average miscalibration is much more pronounced in the knowledge task than in the time series forecasting task (mean index is 5.8 and 1.2, respectively). Although both tasks rely on the interval production method, it is presumably less difficult

[20] A more detailed analysis on the overconfidence measures and their interrelation can be found in Fellner and Krügel (2012).

to give proper confidence intervals in the time series forecasting task because the graphical presentation provides a natural anchor for the lower and the upper bound. Such an anchor is missing in the general knowledge task.[21] This might also explain why we observe greater individual differences in the knowledge task compared to the time series task. In the former, the calibration index ranges from 0 to 10 (SD = 2.4), in the latter, it ranges from 0 to 7 (SD = 1.6).

In the signal-based prediction task, a prediction slope of 0.5 indicates well-calibration, a slope higher than 0.5 reflects overconfidence, and a slope smaller than 0.5 underconfidence. Two different treatments were used to implement the signal-based prediction task: the *Info* treatment and the *No-Info* treatment. Panel A of Table 2 shows that, on average, we observe overconfidence in both treatments. However, with an average prediction slope of 0.87, overconfidence is significantly more pronounced in the *No-Info* treatment compared to the *Info* treatment, where the average prediction slope is 0.65 (Wilcoxon rank sum test: $p < .01$). Recall that subjects in both treatments lacked information about the error variance of the signal, but subjects in the *No-Info* treatment additionally lacked information about the distribution of the outcome variable. Thus, in the *No-Info* treatment, overweighting of signals seems to be a result of two processes: overestimating signal precision (as in the *Info* treatment) and overestimating the variance of the distribution of the outcome variable. Large individual differences are prevalent in both treatments. The prediction slope ranges from 0.36 to 1.02 (SD = 0.16) in the *Info* treatment and from 0.37 to 1.06 (SD = 0.14) in the *No-Info* treatment.

Panel B of Table 2 splits the sample into men and women to investigate possible gender differences in overconfidence since it is often claimed that men are generally more overconfident than women (e.g., Barber and Odean, 2001). Our results, however, do not support this conjecture. In the time series forecasting task and the signal-based prediction task, we do not find significant differences between men and women. In the general knowledge task, we find a significant gender difference in overconfidence (Wilcoxon rank sum test: $p < .05$), but, con-

[21]However, miscalibration indices are found to be highly correlated with interval width scores, calculated by ranking the interval width across participants for each item and summing the ranks for each subject across the ten questions of each task (see Fellner and Krügel, 2012). This indicates that miscalibration in the knowledge and the time series task is in fact due to intervals that are too narrow rather than skewed.

Table 2: Overconfidence measures

A. Descriptive statistics

	Knowledge	Time series	Signal-based predictions Info	Signal-based predictions No-Info
Well-calibrated	1	1	0.5	0.5
No. obs.	168	168	88	80
Mean (SD)	5.8 (2.4)	1.2 (1.6)	0.65 (0.16)	0.87 (0.14)
Min	0	0	0.36	0.36
Max	10	7	1.02	1.06

B. Gender differences

	Male		Female		p-value
	Obs.	Mean (SD)	Obs.	Mean (SD)	
Knowledge	83	5.43 (2.55)	85	6.18 (2.18)	0.03**
Time series	83	1.05 (1.49)	85	1.26 (1.70)	0.52
Signal-based (Info)	43	0.62 (0.16)	45	0.67 (0.14)	0.15
Signal-based (No-Info)	40	0.87 (0.16)	40	0.87 (0.12)	0.31

C. Pairwise correlations

	Knowledge	Time series	Obs.
Knowledge	1		168
Time series	0.50 (0.000)***	1	168
Signal-based (Info)	0.17 (0.109)	0.31 (0.003)***	88
Signal-based (No-Info)	0.09 (0.409)	-0.05 (0.644)	80

Notes: p-values in Panel B refer to the results of Wilcoxon rank sum tests.
*** and ** indicate significance at the 1% and 5% margin, respectively.

trary to the frequent claim, our results show that women are more overconfident than men in this task (mean calibration scores are 6.18 and 5.43, respectively). We thus reject Hypothesis 5, and thereby also the conjecture that men trade more than women because they are more overconfident in judgment tasks (Hypothesis 6).[22]

Panel C of Table 2 displays all pairwise correlations between the overconfidence measures. If the implicit assumption of previous empirical tests of the overconfidence hypothesis holds, all three tasks should uncover the same underlying judgmental bias. In this case, a significant and

[22]In our multivariate analysis in section 4.2.2, we have also investigated whether overconfidence affects trading activity differently for men and women. However, none of the interactions between gender and overconfidence were significant, which supports the result that gender differences in trading are not due to any gender differences in overconfidence.

positive correlation between all overconfidence measures is to be expected. Panel C shows that this is not the case. Individuals who tend to be miscalibrated in the general knowledge task tend to be miscalibrated in the time series forecasting task as well ($\rho = .5$, $p < .01$), but the correlation between miscalibration and overweighting of signals is far less pronounced. Only for the *Info* treatment do we find a positive and significant correlation between miscalibration in time series forecasts and overweighting of signals ($\rho = .3$, $p < .01$).[23] Miscalibration in general knowledge and overweighting of signals is never significantly correlated, neither in the *Info* treatment nor in the *No-Info* treatment. Thus, miscalibration in general knowledge and time series forecasting seems to capture a common judgmental bias, while the (mis-)perception of signal reliability appears to be a distinct phenomenon. This lends support to the supposition that previous empirical tests of the overconfidence hypothesis have conceptualized overconfidence in a way that is distinct from the modeled bias in the theoretical literature.

4.2 Individual trading activity

4.2.1 Descriptive and bivariate statistics

Table 3 gives an overview of individual trading activity which is measured by the average number of concluded trades across all 10 trading rounds.[24] Panel A shows that, on average, subjects concluded 10.4 trades per round, where the overwhelming majority of trades took place in the double auction (7.42). In the call market, average individual trading activity was less than half of that in the double auction (2.95). Panel A also shows large individual differences in trading activity. Taking both market phases together, the subject with the lowest trading activity concluded 0.9 trades per round and the subject with the highest trading activity concluded 39.5 trades per round (SD = 7.28). In the double auction, individual trading activity ranged from 0 to 32.2 (SD = 5.69) while in the call market, it ranged from 0 to 15.3 (SD = 2.68).

[23] When we control for other influencing factors such as interval width scores, the correlation coefficient almost halves and is only marginally significant.

[24] Averaging across rounds is common practice in experimental financial studies to filter out some noise in the data (e.g., Biais et al., 2005). However, it should be noted that this assumes independence of trading rounds. In our experiment, all subjects remained within one market and thus shared a common history. But since endowments were re-initialized and only one trading round was paid out, a certain degree of independence was established.

Panel B of Table 3 demonstrates a pronounced gender difference in individual trading activity. On average, men concluded significantly more trades per round than women (Wilcoxon rank sum test: $p < .01$). The observed gender difference is also substantial as the average trading activity of men was approximately 34% higher than that of women. Moreover, it seems to be independent of the market type. Men traded significantly more than women in the call market as well as in the continuous double auction (Wilcoxon rank sum tests: $p < .01$ and $p < .05$, respectively). The last line of Panel B points to a possible explanation of the gender difference in trading activity. According to our risk measure, men are significantly less risk averse than women (3.23 vs 2.67, Wilcoxon rank sum test: $p < .05$).[25]

Panel C of Table 3 provides a first glimpse at the question whether judgmental overconfidence affects trading volume. For each overconfidence measure, we divided our subjects into two groups based on a median-split and calculated the average trading activity in each group. In line with prior studies, we do not find a significant difference in trading activity when the median-split is based on either of the two miscalibration measures. The same applies to the median-split of the prediction slopes in the *Info* treatment. In contrast, when the comparison between the two groups is based on the prediction slopes in the *No-Info* treatment, we find a marginally significant difference in individual trading activity (Wilcoxon rank sum test: $p = .07$). Subjects with an overweighting bias above the median concluded approximately 22% more trades than subjects with an overweighting bias below the median.

The last line of Panel C suggests that trading activity also depends on risk attitudes. Subjects with a risk index above the median, indicating lower degrees of risk aversion, traded significantly more than subjects with a risk index below the median ($p = 0.021$). Together with the evidence that men are, on average, less risk averse than women, this strengthens the supposition that the gender difference in trading activity might be driven by different degrees of risk aversion.

[25] For 12 of the 168 participants, no risk index could be calculated due to non-monotonous switching behavior between lotteries. Although not rational, non-monotonous switching behavior is always observed for a small group of individuals in this lottery choice task. In addition, three further subjects were not assigned a risk index because they violated the dominant choice in lottery number one or eleven. The total number of unclassifiable subjects is well in line with findings of other authors who used this method.

Table 3: Trading activity

A. Descriptive statistics

	Trades total	Continuous market	Call market
No. obs.	168	168	168
Mean (SD)	10.4 (7.28)	7.42 (5.69)	2.95 (2.68)
Min	0.9	0	0
Max	39.5	32.2	15.3

B. Gender differences in trading activity and risk attitude

	Male		Female		p-Value
	Obs.	Mean (SD)	Obs.	Mean (SD)	
Trades total	83	11.9 (7.96)	85	8.89 (6.24)	0.009***
Cont. market	83	8.27 (6.17)	85	6.60 (5.06)	0.046**
Call market	83	3.63 (2.98)	85	2.29 (2.16)	0.001***
Risk attitude	78	3.23 (1.50)	75	2.67 (1.38)	0.024**

C. Group differences in total trading activity, based on median-splits

	Score<Median		Score>Median		p-Value
	Obs.	Mean (SD)	Obs.	Mean (SD)	
Knowledge	70	9.63 (6.56)	72	11.32 (7.92)	0.178
Time series	80	9.89 (6.49)	45	11.93 (8.65)	0.287
Signal-based (Info)	44	9.74 (6.82)	44	11.00 (7.94)	0.385
Signal-based (No-Info)	40	9.37 (7.54)	40	11.39 (6.80)	0.073*
Risk attitude	65	9.39 (8.05)	54	11.90 (7.92)	0.021**

Notes: p-values refer to the results of Wilcoxon rank sum tests.
***, ** and * indicate significance at the 1%, 5% and 10% margin, respectively.

4.2.2 Multivariate analysis

We now turn to a multivariate analysis by regressing individual trading activity on several overconfidence measures, risk attitude as well as some available demographic characteristics of participants. Table 4 displays the coefficients and standard errors from mixed-effects regressions with a random group intercept.[26] The dependent variable, individual trading activity, is the log-transformed average number of concluded trades for each individual across all 10 trading

[26]The indicated levels of significance in the table are based on the t-distribution. However, p-values based on the posterior density distribution of the parameters via Markov chain Monte Carlo simulation (n=10,000) are virtually identical and lead to the same conclusions.

rounds.[27]

Table 4: Regression analysis for total trades

Dependent variable: Log of average total number of trades
Method: Mixed-effects ML estimation with group random effects

Indep. var.	(1) All data	(2) All data	(3) All data	(4) Info	(5) No-Info	(6) Info	(7) No-Info
OC Knowledge		-0.006				0.021	-0.047
		(0.026)				(0.043)	(0.034)
OC Time series			-0.036			0.004	-0.087
			(0.043)			(0.069)	(0.057)
OC Signal rel.				-0.677	1.017**	-0.727	1.193**
				(0.597)	(0.518)	(0.622)	(0.499)
Male	0.618**	0.610**	0.587**	0.504	0.945***	0.535	0.851**
	(0.269)	(0.270)	(0.269)	(0.361)	(0.367)	(0.369)	(0.351)
Risk	0.170***	0.171***	0.169***	0.259***	0.137*	0.249***	0.118*
	(0.061)	(0.061)	(0.061)	(0.092)	(0.074)	(0.094)	(0.071)
Male*Risk	-0.145*	-0.143*	-0.135*	-0.240**	-0.135	-0.241**	-0.095
	(0.082)	(0.082)	(0.082)	(0.115)	(0.108)	(0.119)	(0.103)
Age	0.034	0.034	0.035	0.042	0.038	0.045	0.035
	(0.029)	(0.029)	(0.029)	(0.045)	(0.039)	(0.045)	(0.037)
Semester	-0.058**	-0.058**	-0.062***	-0.047	-0.083***	-0.044	-0.098***
	(0.023)	(0.023)	(0.023)	(0.038)	(0.027)	(0.038)	(0.026)
Constant	1.051*	1.101*	1.088*	1.199	0.135	1.051	0.561
	(0.621)	(0.652)	(0.620)	(0.863)	(1.003)	(0.915)	(0.961)
σ_u	0.222***	0.224***	0.232***	0.309***	0.189***	0.307***	0.240***
	(0.084)	(0.084)	(0.084)	(0.114)	(0.117)	(0.114)	(0.105)
σ_i	0.703***	0.703***	0.700***	0.696***	0.618***	0.695***	0.582***
	(0.043)	(0.043)	(0.043)	(0.060)	(0.055)	(0.060)	(0.052)
No. of obs.	153	153	153	79	74	79	74
No. of groups	21	21	21	11	10	11	10

Notes: ***, ** and * indicate significance at the 1%, 5%, and 10% margin, respectively.

Specification 1 presents the estimates of the basic model where no overconfidence measure is included. In each regression reported in subsequent columns, we add one of the overconfidence measures to the basic specification. Specifications 6 and 7 contain the estimates of the full model where all overconfidence measures are included. Note that the regressions presented in

[27]We log-transformed the average number of concluded trades to avoid problems like non-normality. We used a log transformation for the sake of a straightforward interpretation of coefficients. However, several robustness checks, e.g., a mixed-effects negative binomial regression, a Box-Cox transformation for zero skewness of the trade variable, and separate regressions for the first and second half of trading rounds confirm the results.

specifications 1 to 3 are based on the joint data, while specifications 4 to 7 are based on the separated data for the *Info* and *No-Info* treatment of the signal-based prediction task.

Results of the basic model in specification 1 show that individual trading activity depends on both the gender of participants and their risk attitude. However, the (marginally) significant interaction between the two variables indicates that the effect of gender and risk attitude on trading is more complex than expected. For women we find that the influence of risk attitude on trading activity is substantial, as indicated by the significant coefficient on the risk variable: the less risk averse a woman in the binary lottery choice task, the more trades she concludes. For men, on the other hand, the risk attitude seems to play only a minor role. In fact, our risk measure is not significantly related to men's trading activity.[28] Consequently, we find evidence in favor of Hypothesis 3 for the subsample of female traders but not for the subsample of male traders.

Regarding gender differences in trading volume, our results indicate that men trade significantly more than women, though only at high levels of risk aversion. The gender gap narrows as risk aversion decreases. In fact, at levels approaching risk neutrality, men trade equally as much as women.[29] We conclude that gender differences in trading do not disappear, once we control for the risk attitude of our subjects, but they are only present at high levels of risk aversion and become less pronounced as risk aversion decreases. Thus, Hypothesis 4 has to be rejected. These results might explain the mixed evidence on gender differences in trading activity in prior studies. Whereas some studies find that men trade more than women (e.g., Barber and Odean, 2001, Biais et al., 2005, Grinblatt and Keloharju, 2009), others do not find a gender effect on trading (e.g., Glaser and Weber, 2007, Grinblatt and Keloharju, 2001, Deaves et al., 2009). Our results suggest that differences in average levels of risk aversion across studies might reconcile existing contradictory evidence.

The basic model also shows that the semester variable, which might be interpreted as progress in studies, has a negative and significant effect on trading activity, whereas age seems

[28]This becomes obvious by testing the linear combination of Risk and Male*Risk against zero ($p = .64$).

[29]Testing the linear combination of Male and Male*Risk at a risk index of 5 reveals no significant difference from zero: $p = .61$.

to be unrelated to trading volume in our sample.[30] Also note that the results on gender, risk attitude, and semester are robust to the inclusion of measures of miscalibration as long as the joint data is used (specifications 2 and 3). For the subset of the data in the *Info* treatment, semester becomes insignificant (specifications 4 and 6) while for the subset of the data in the *No-Info* treatment, the interaction between gender and risk attitude becomes insignificant (specifications 5 and 7).

We now turn to the main question of interest, i.e., the influence of overconfidence on trading. Specifications 2 and 3 of Table 4 reveal that miscalibration based on the knowledge and time series forecasting tasks are unrelated to individual trading activity. This is in line with the findings of Biais et al. (2005) and Glaser and Weber (2007), lending support to Hypothesis 1. Specifications 4 and 5 show the results for the two treatments of the signal-based prediction task. In the *Info* treatment, the prediction slope appears to be unrelated to individual trading activity. This is surprising since the overconfidence measure based on this treatment corresponds most closely to the modeled judgmental bias in the theoretical literature. However, the estimation results in specification 5 show that the link between individual trading activity and the prediction slope in the *No-Info* treatment is positive and significant ($p = 0.05$). These results are confirmed by the full regression model where all overconfidence measures are included at once (specifications 6 and 7). Whereas none of the miscalibration measures nor the prediction slope based on the *Info* treatment are significantly related to trading volume, the prediction slope based on the *No-Info* treatment remains significantly linked to individual trading activity in the expected direction ($p = 0.022$). The effect is also substantial: an increase of 0.14 points in the overconfidence score (which corresponds to one standard deviation) is associated with approximately 14% to 17% more trade (specifications 5 and 7, respectively).

To make use of the data in the independent prior call market and the subsequent continuous double auction separately, Table 5 repeats the above described analysis for each of the two markets. In both markets, the dependent variable is the average number of trades throughout

[30]Excluding age from the regressions does not change any of the results. To show that the effect of the semester variable is not driven by subjects' age, we decided to present the regressions where the age variable is included.

the 10 rounds. Some subjects never traded in one of the two markets so that a simple log transformation of the dependent variable would eliminate these subjects from the analysis. We therefore decided to transform the average number of trades in the call and double auction markets by first adding a constant equal to the square of the third quartile divided by the first quartile to all data (see Stahel, 2002), followed by a Box-Cox transformation to obtain a distribution with zero skewness.[31] The resulting variable is normally-distributed but implies that the quantitative interpretation of the regression coefficients is less straightforward than the results presented in Table 4.[32]

The results on risk attitude and gender are similar to the previous analysis. Higher risk tolerance is found to significantly increase trading volume in all markets except in the call market of the *No-Info* treatment. Men tend to trade more than women at high levels of risk aversion in both market types, but the difference is only (marginally) significant in the *No-Info* treatment. Male traders also appear to be less sensitive to risk preferences than female traders, but the effect is statistically weak as it is only marginally significant in the continuous double auction of the *Info* treatment. The progress in studies (i.e., semester) is negatively related to trading volume in three of the four regressions.

Turning again to the main question of the link between overconfidence and trading activity, we confirm the aggregate findings that miscalibration in general knowledge and time series predictions appears to be unrelated to trading in either market or subsample. Also, in the call market, we find no relation between overconfidence in the signal-based prediction task and trading volume. In the continuous market, however, overconfidence in signal perception predicts individual trading activity.[33] In the *No-Info* treatment (specification 4), higher overconfidence

[31]Alternatively, applying a simple log transformation of the data after adding a constant of half of the smallest positive value leaves the qualitative results of the subsequent regressions unchanged.

[32]We ran several robustness checks for our results like negative binomial regression models on the aggregated data as well as a Poisson and a Hurdle model on the disaggregated data. They are not reported here for the sake of brevity, but the results reported in Table 5 are robust throughout.

[33]Note that Biais et al. (2005) found that psychological characteristics of individuals have a higher impact in a call market than in a double auction. We find the opposite: overconfidence in signal-based predictions is significantly related to trading activity in the double auction but not in the call market. However, the variable of interest in the study by Biais et al. (2005) was performance (i.e., earnings), whereas in our study it is individual trading volume. Recall also that most of the trading in our experimental asset market took place in the continuous double auction. In the call market, subjects tended to be very cautious and preferred to state wide bid-ask spreads. As a consequence, in many trading rounds the bids and asks of different subjects did

Table 5: Regression analysis for call and continuous markets

Dependent variable: Box-Cox transformed average number of trades on each market
Method: Mixed-effects ML estimation with group random effects

	Call market		Double auction	
Independent variable	(1) Info	(2) No-Info	(3) Info	(4) No-Info
OC Knowledge	0.072	-0.039	0.023	-0.070
	(0.081)	(0.073)	(0.065)	(0.051)
OC Time series	0.031	-0.034	0.019	-0.117
	(0.127)	(0.123)	(0.105)	(0.086)
OC Signal rel.	0.573	0.310	-1.818*	1.944***
	(1.199)	(1.087)	(0.944)	(0.755)
Male	1.128	1.388*	0.499	1.223**
	(0.718)	(0.765)	(0.559)	(0.531)
Risk	0.389**	0.062	0.284**	0.177*
	(0.180)	(0.154)	(0.143)	(0.107)
Male*Risk	-0.288	-0.038	-0.299*	-0.173
	(0.231)	(0.226)	(0.180)	(0.157)
Age	0.074	0.042	0.062	0.037
	(0.085)	(0.080)	(0.069)	(0.056)
Semester	-0.020	-0.127**	-0.095*	-0.135***
	(0.074)	(0.058)	(0.058)	(0.040)
Constant	-2.654	-0.193	1.840	0.257
	(1.714)	(2.089)	(1.399)	(1.452)
σ_u	0.142	0.642	0.552*	0.403**
	(0.461)	(0.238)	(0.179)	(0.157)
σ_e	1.367***	1.259**	1.052	0.877
	(0.118)	(0.113)	(0.090)	(0.078)
No. of obs.	79	74	79	74
No. of groups	11	10	11	10

Notes: ***, ** and * indicate significance at the 1%, 5%, and 10% margin.

is significantly associated with higher trading volume, as found before in the regressions with aggregated trading volume. In the *Info* treatment (specification 3), the link between overconfidence and trading activity appears to be marginally significant, but the direction of the relation is reversed: higher overconfidence leads to less trading. Whereas the finding in the *No-Info* treatment is in line with the predictions of the overconfidence models, the result in the *Info*

not overlap such that no trading took place at all. This general reluctance of our subjects to trade in the call market might be a reason for the weak effects in the regressions for this market type.

treatment contradicts these models.

Recall that the only difference between the signal-based prediction task in the *Info* and *No-Info* treatment was the prior information about the underlying distribution of the outcome variable. Whereas subjects in the *Info* treatment were informed about the distribution, subjects in the *No-Info* treatment were not. Thus, in the *No-Info* treatment, subjects had to learn about the precision of the signal (just as in the *Info* treatment) as well as about the variability of the outcome distribution. Since the overweighting parameter, which is based on the *No-Info* treatment but not on the *Info* treatment, is positively related to trading activity, it seems that misperceiving the variability of the outcome distribution is the crucial behavioral factor that increases trading. Interestingly, Odean (1998) shows in his model that underweighting common priors (which corresponds to overestimating the outcome variability) increases trading volume just as overestimating signal precision. However, he deems the latter bias more important. Our results suggest the opposite.

Nevertheless, the weakly negative relationship between individual trading activity and overconfidence based on the *Info* treatment remains peculiar, and we can only speculate about possible causes. Since subjects in this treatment received all information about the underlying distribution of the outcome variable, some might have anchored their predictions in the expected value of the outcome distribution. Such anchoring and, subsequently, insufficient adjustment to the signal will necessarily make the prediction slope flatter, even if the precision of the signals is overestimated. Thus, the prediction slope of some subjects might indicate only little overweighting of private information, even though they highly overestimated the precision of their signals. We would expect that such a downward bias of the prediction slope is the more severe, the more the precision of the signal is overestimated, because more adjustment toward the signal would have been necessary. As it cannot be ruled out that the measurement of overconfidence is confounded by an additional anchoring-and-adjustment heuristic, the interpretation of the negative relation between the prediction slope in the *Info* treatment and individual trading activity requires caution. In terms of extrapolating the results to market environments outside the lab, we believe, however, that the *No-Info* treatment captures the more realistic setting since real trading environments are usually devoid of precise information about the distributional properties of the asset value.

5 Summary and Conclusion

Economic models explain the excessively high trading activity observed in financial markets by overconfident traders. Overconfidence in these models is reflected by traders who overestimate the precision of private signals and, consequently, overweight this private information when updating their beliefs. Previous empirical tests have by and large found no evidence for the modeled link between overconfidence and trading activity. We argue that it might be too early to discard (judgmental) overconfidence as a reason for excessive trading due to a discrepancy between modeling and measuring overconfidence: all previous studies rely on classical measures of miscalibration based on general knowledge and time series prediction tasks. It has been shown, however, that miscalibration in these tasks is unrelated to the perception of signal reliability.

We conducted an experimental asset market, consisting of a call and a continuous market phase, to examine the relation of trading activity and overconfidence once again. To assess overconfidence, we obtained three measures: the two classical ones, miscalibration in knowledge questions and miscalibration in time series predictions, and a third measure that captures overconfidence in signal-based predictions. We employed two treatments for the signal-based prediction task. In the *Info* treatment, overconfidence may arise from underestimating the noise of the signal, while in the *No-Info* treatment, overconfidence could additionally arise from overestimating the variance of the outcome variable. According to trading models (e.g., Odean, 1998), both effects should lead to higher trading volume. In contrast to previous studies, we additionally controlled for risk attitude and analyzed gender effects.

We find that, on average, men trade more than women, but the gender difference in trading activity becomes less pronounced for lower degrees of risk aversion. The latter finding is mainly due to the fact that the trading activity of women in our sample is more sensitive to risk preferences than the trading activity of men. Our main analysis replicates the finding that miscalibration is indeed not related to trading activity. However, overconfidence in signal-based predictions is associated with trading activity, particularly in the continuous double auction market. Yet the predicted positive relation can only be found for the *No-Info* treatment: individuals with higher overconfidence in this treatment tend to trade more. In the *Info* treatment, the relation appears to be reversed. However, the results in the *Info* treatment should be treated with caution since our overconfidence measure in this treatment might be confounded by an additional anchoring-and-adjustment bias, which operates in the opposite direction.

The conclusions to be drawn from these findings are twofold: First, and most important, we demonstrate that the hitherto lack of support for overconfidence trading models may be rooted in the improper empirical conceptualization of overconfidence. Miscalibration in general knowledge and time series prediction tasks is unrelated to overconfidence in the domain of signal perception and does not predict trading activity. The vast literature in psychology has already established that overconfidence is less robust across domains than initially assumed (e.g., Klayman et al., 1999). Therefore it is essential to consider the relevant domain to measure overconfidence in order to examine its possible ramifications on market outcomes.

Second, even when using an empirical measure of overconfidence that is as close as possible to the theoretical models, the evidence is not unambiguously supportive of the theory. Our findings indicate that it may be the additional effect of overestimating the outcome variability in the *No-Info* treatment (i.e., the underweighting of common priors) which is behaviorally relevant for excessive trading. We interpret this result as tentative support for the overconfidence trading models, albeit the causal relationship is not the one most emphasized in these models.

References

Barber, B., Odean, T., 2001. Boys will be boys: Gender, overconfidence, and common stock investment. Quarterly Journal of Economics 116 (1), 261–292.

Benos, A., 1998. Aggressiveness and survival of overconfident traders. Journal of Financial Markets 1, 353–383.

Biais, B., Hilton, D., Mazurier, K., 2005. Judgmental overconfidence, self-monitoring, and trading performance in an experimental financial market. The Review of Economic Studies 72 (2), 287–312.

Caballé, J., Sákovics, J., 2003. Speculating against an overconfident market. Journal of Financial Markets 6 (2), 199–225.

Chuang, W., Lee, B., 2006. An empirical evaluation of the overconfidence hypothesis. Journal of Banking & Finance 30 (9), 2489–2515.

Czaczkes, B., Ganzach, Y., 1996. The natural selection of prediction heuristics: Anchoring and adjustment versus representativeness. Journal of Behavioral Decision Making 9 (2), 125–139.

De Bondt, W., Thaler, R., 1995. Financial decision-making in markets and firms: A behavioral perspective. Handbooks in Operations Research and Management Science 9, 385–410.

Deaves, R., Lüders, E., Luo, G., 2009. An experimental test of the impact of overconfidence and gender on trading activity. Review of Finance 13, 555–575.

Eckel, C. C., Grossman, P. J., 2008. Men, women and risk aversion: Experimental evidence. In: Handbook of Experimental Economics Results. Vol. 1. North-Holland, pp. 1061–1073.

Fellner, G., Krügel, S., 2012. Judgmental Overconfidence: Three Measures, One Bias? Journal of Economic Psychology 33 (1), 142 – 154.

Fellner, G., Maciejovsky, B., 2007. Risk attitude and market behavior: Evidence from experimental asset markets. Journal of Economic Psychology 28 (3), 338–350.

Fischbacher, U., 2007. Zurich toolbox for ready-made economic experiments. Experimental Economics 10 (2), 171–178.

Ganzach, Y., 1993. Predictor representation and prediction strategies. Organizational Behavior and Human Decision Processes 56, 190–212.

Ganzach, Y., 1994. Feedback representation and prediction strategies. Organizational Behavior and Human Decision Processes 59, 391–409.

Glaser, M., Weber, M., 2007. Overconfidence and trading volume. The Geneva Risk and Insurance Review 32 (1), 1–36.

Greiner, B., 2004. An online recruitment system for economic experiments. In: Kremer, K., Macho, V. (Eds.), Forschung und wissenschaftliches Rechnen, GWDG Bericht 63. Gesellschaft für Wissenschaftliche Datenverarbeitung, Göttingen, pp. 79–83.

Grinblatt, M., Keloharju, M., 2001. What makes investors trade? The Journal of Finance 56 (2), 589–616.

Grinblatt, M., Keloharju, M., 2009. Sensation seeking, overconfidence, and trading activity. The Journal of Finance 64 (2), 549–578.

Hilton, D., Régner, I., Cabantous, L., Charalambides, L., Vautier, S., 2011. Do positive illusions predict overconfidence in judgment? A test using interval production and probability evaluation measures of miscalibration. Journal of Behavioral Decision Making 24 (2), 117–139.

Holt, C. A., Laury, S. K., 2002. Risk aversion and incentive effects. American Economic Review 92, 1644–1655.

Klayman, J., Soll, J., González-Vallejo, C., Barlas, S., 1999. Overconfidence: It depends on how, what, and whom you ask. Organizational Behavior and Human Decision Processes 79 (3), 216–247.

Kogan, S., 2009. Distinguishing the effect of overconfidence from rational best-response on information aggregation. Review of Financial Studies 22 (5), 1889.

Kyle, A., Wang, F., 1997. Speculation duopoly with agreement to disagree: Can overconfidence survive the market test? Journal of Finance 52, 2073–2090.

Lawrence, M., O'Connor, M., 1992. Exploring judgmental forecasting. International Journal of Forecasting 8 (1), 15–26.

Lawrence, M., O'Connor, M., 1993. Scale, variability, and the calibration of judgmental prediction intervals. Organizational behavior and human decision processes 56 (3), 441–458.

Lichtenstein, S., Fischhoff, B., Phillips, L., 1982. Calibration and probabilities: The state of the art to l980. Judgment under uncertainty: Heuristics and biases, 306–334.

Lundeberg, M. A., Fox, P. W., Punccohar, J., 1994. Highly confident but wrong: Gender differences and similarities in confidence judgments. Journal of Educational Psychology 86, 114–121.

Moore, D., Healy, P., 2008. The trouble with overconfidence. Psychological Review 115 (2), 502–517.

Odean, T., 1998. Volume, volatility, price and profit: when all traders are above average. Journal of Finance 53, 1887–1934.

Russo, J., Schoemaker, P., 1992. Managing overconfidence. Sloan Management Review 33 (2), 7–17.

Stahel, W. A., 2002. Statistische Datenanalyse. Eine Einführung für Naturwissenschaftler. Vieweg, Braunschweig.

Statman, M., Thorley, S., Vorkink, K., 2006. Investor overconfidence and trading volume. Review of Financial Studies 19 (4), 1531.

Weber, E. U., Blais, A.-R., Betz, N. E., 2002. A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. Journal of Behavioral Decision Making 15, 263–290.

A Appendix

A.1 General knowledge questions

The following ten questions were used in the general knowledge task. Correct answers are in parentheses.

1. What is the length of the River Nile in km? (6,671 km)
2. How many states are currently *(Nov. 2009)* members of OPEC? (12)
3. What is the average diameter of the Moon in km? (3,745 km)
4. What was the number of inhabitants of Australia in 2008 (in Mill.)? (21.374 Mill.)
5. What is the number of passenger airports in Germany? (38)
6. What was the number of patent applications in Germany in 2008? (62,417)
7. What is the size of France in km^2? (674,843 km^2)
8. What is the air distance between London and Tokyo in km? (9,581 km)
9. When was the novel Robinson Crusoe by Daniel Defoe first published? (1719)
10. When was the zip fastener patent registered? (1893)

A.2 Time series forecasting task

The following ten time series were used in the time series forecasting task and were pre-generated using an autoregressive, moving average process with one MA and one AR term.

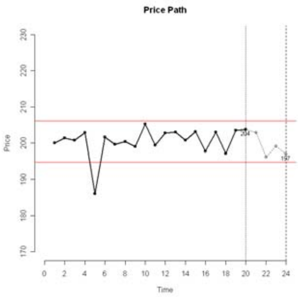

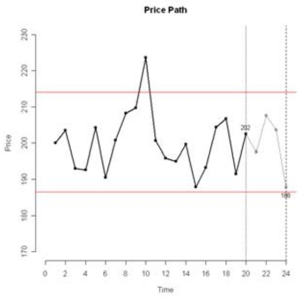

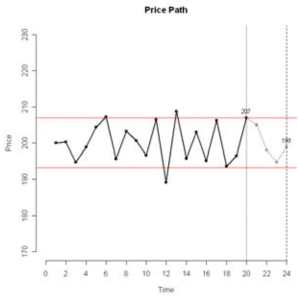

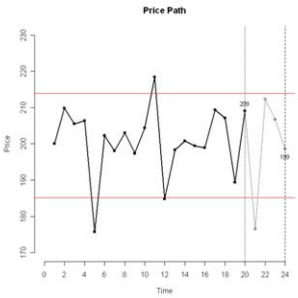

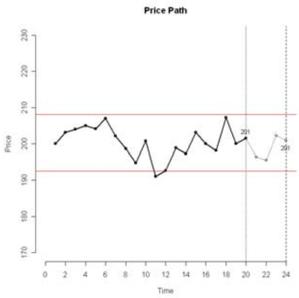

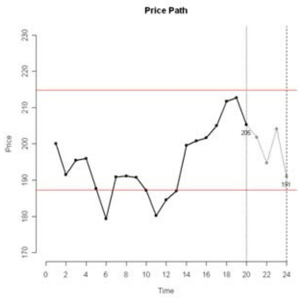

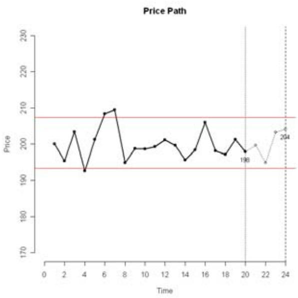

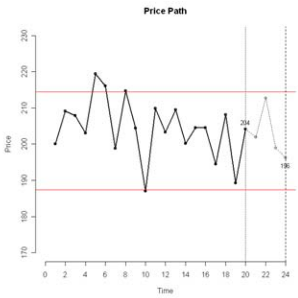

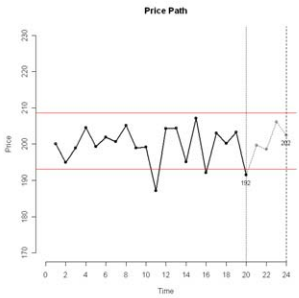

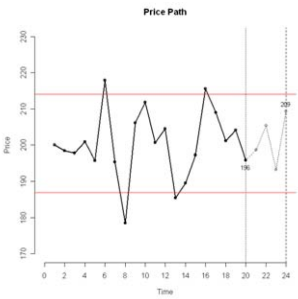

A.3 Instructions

Original instructions translated from German

Thank you for participating in this experiment. Please, do not talk to other participants from now on. If you have any questions, please raise your arm and one of the experimenters will answer your question at your cubicle. Please do not forget to switch off your mobile phone.

General information

For showing up on time at the laboratory, you receive 2.50 €. Depending on your decisions, you can earn additional money in the experiment. The amount of 2.50 € and the additional money you earn will be paid out to you in cash at the end of the experiment. The payout will be done privately for each participant, so no-one else will know the amount you earned. All decisions are anonymous and cannot be traced to the name of the participant.

The experiment is divided into 4 phases. At the beginning of each phase you obtain new instructions. To understand the procedure of each phase, please read the instructions carefully!

Phase 1

Phase 1 consists of two parts. In each part, there are ten questions; that means 20 questions in total. The ten questions of part one are almanac questions. All questions in part one are similar to the following one:

Example: How many member states does the European Union have?

It is not expected that you know the exact answer to each question. To answer each of the 10 questions, you have to state a range of numbers so that you are 90% certain that the true answer lies within this range. To be 90% certain roughly means that out of the 10 questions, the correct answer should lie outside your stated range for only one of the question. Thus, for each question you will have to state two numbers:

(a) An **upper** bound so that you are 95% certain that the true answer lies *below* this number.

(b) A **lower** bound so that you are 95% certain that the right answer lies *above* this number.

Once you have answered all 10 questions of part one, you will automatically start with part two. The 10 questions of part two are concerned with time series predictions. A time series is a chronological sequence of data like, for instance, data of stock prices. You can find an example for such a time series at the end of these instructions. All 10 time series in this experiment were computer-generated and reflect the price movement of a **hypothetical** stock. None of the 10 time series is based on the price path of an actual stock so that it is not possible to recognize price patterns of real stocks.

All 10 time series start with an initial price of 200 € in period 1. Starting from this value, a price for 23 subsequent periods was computer-generated for all 10 time series. In total, each time series consists of the price path of a hypothetical stock over 24 periods. Yet, you will see only the prices of the first **20** periods. It is then your task to predict the price of the stock in period **24** based on the prices of the first 20 periods. Just like in part 1, it is not expected that you can predict the price exactly. Like in part 1, you are asked to give a price range so that you are 90% certain that the actual price in period 24 lies within this range. Being 90% certain roughly means (like in part 1) that out of 10 time series the actual price in period 24 should lie outside the stated price range in only one case. Thus, for each time series you will have to state two numbers:

(a) An **upper** price limit so that you are 95% certain that the actual price in period 24 lies *below* this number.

(b) A **lower** price limit so that you are 95% sure that the actual price in period 24 lies *above* this number.

Payoff for phase 1:
When you have answered all 20 questions of phase 1, 3.00 € will be added automatically to your earnings in the experiment.

Example for a time series:

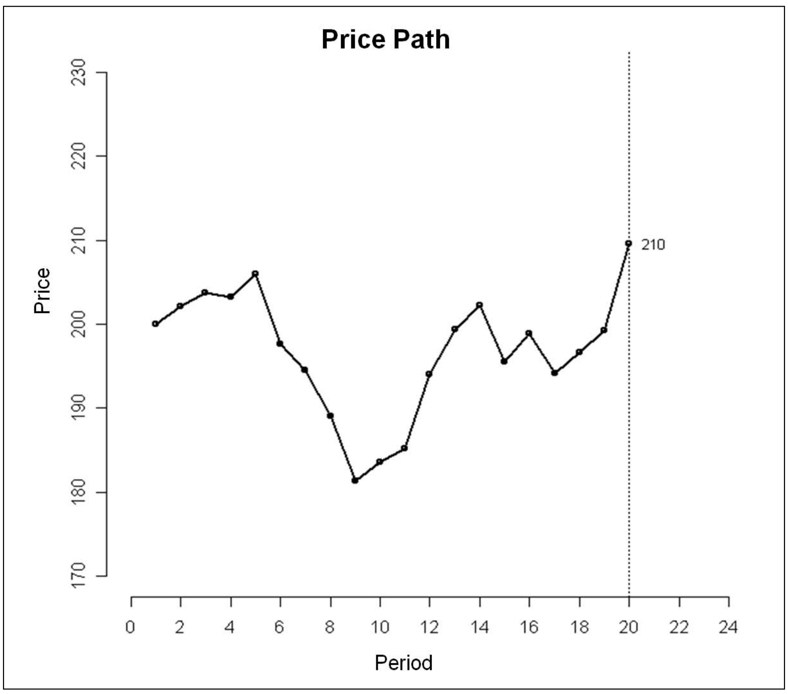

Phase 2

In this phase, you have to choose either option X or option Y in 11 different cases. All 11 cases will be presented in a list at once on your screen. Each of the two options has 2 possible monetary outcomes **(all values in Euro)**, one high outcome and one low outcome, each of which will be paid out with certain probabilities. Whereas the two possible monetary outcomes of both options remain the same in all 11 cases, the probabilities with which they will be paid out alter.

Options X and options Y will be presented to you in the following way:

Lottery Task

Please choose either option X or option Y in all 11 cases. At the end of the experiment, one of the 11 cases will be selected at random for your payment.

All values in Euro!

	Option X	Option Y		
1.	with 0/10 outcome of 2, with 10/10 outcome of 1.60	with 0/10 outcome of 3.85, with 10/10 outcome of 0.10	X	Y
2.	with 1/10 outcome of 2, with 9/10 outcome of 1.60	with 1/10 outcome of 3.85, with 9/10 outcome of 0.10	X	Y
3.	with 2/10 outcome of 2, with 8/10 outcome of 1.60	with 2/10 outcome of 3.85, with 8/10 outcome of 0.10	X	Y
4.	with 3/10 outcome of 2, with 7/10 outcome of 1.60	with 3/10 outcome of 3.85, with 7/10 outcome of 0.10	X	Y
5.	with 4/10 outcome of 2, with 6/10 outcome of 1.60	with 4/10 outcome of 3.85, with 6/10 outcome of 0.10	X	Y
6.	with 5/10 outcome of 2, with 5/10 outcome of 1.60	with 5/10 outcome of 3.85, with 5/10 outcome of 0.10	X	Y
7.	with 6/10 outcome of 2, with 4/10 outcome of 1.60	with 6/10 outcome of 3.85, with 4/10 outcome of 0.10	X	Y
8.	with 7/10 outcome of 2, with 3/10 outcome of 1.60	with 7/10 outcome of 3.85, with 3/10 outcome of 0.10	X	Y
9.	with 8/10 outcome of 2, with 2/10 outcome of 1.60	with 8/10 outcome of 3.85, with 2/10 outcome of 0.10	X	Y
10.	with 9/10 outcome of 2, with 1/10 outcome of 1.60	with 9/10 outcome of 3.85, with 1/10 outcome of 0.10	X	Y
11.	with 10/10 outcome of 2, with 0/10 outcome of 1.60	with 10/10 outcome of 3.85, with 0/10 outcome of 0.10	X	Y

O K

For the second case, for instance, this means:
Option X pays out either 2.00 € with probability 1/10 or 1.60 € with probability 9/10.
Option Y pays out either 3.85 € with probability 1/10 or 0.10 € with probability 9/10.

On the right side of the screen you have to tick your preferred option for each of the 11 cases. The left circle selects option X, the right circle option Y.

Payoff for phase 2:
At the end of the experiment (after phase 4), one of the 11 cases will be selected at random by the computer. All cases are equally likely.

For the selected case, the computer then randomly picks the high or the low monetary outcome of your chosen option according to the effective probabilities.

Phase 3

In this phase of the experiment, you have to predict the value of a variable X. As an aid for this task, you will receive a signal that is related to the variable X. This means that the signal hints at the value of the variable X, but it is subject to an error. The signal can be understood as an indicator for the value of the variable X. The relation between the variable X and the signal is positive. That means, on average, the higher the signal, the higher the value of the variable X.
In several rounds, you have to predict the value of the variable X based on the signal. It is nearly impossible to predict the value of the variable X exactly. Therefore, you are asked to give a prediction that is as close as possible to the actual value of the variable X.

Procedure:
There are 60 rounds in total which means that you make 60 predictions. In each round, you have to predict a new value of the variable X based on a new signal. The value of the variable X and the value of the signal in each round are independent of previous values of the variable X and previous values of the signal. All values are independent across rounds.
At the beginning of each round, you receive the signal and based on this information you have to predict the value of the variable X in this round. After you submitted your prediction, the true value of the variable X in the corresponding round is displayed on screen. Additionally, on the same screen, you will again see the signal of this round, your prediction and the absolute deviation of your prediction from the actual value of the variable X.
One round lasts for about 30 seconds. That means, for about 15 seconds, you see the signal and you make your prediction. After clicking the button "OK," the actual value of the variable X in this round and the additional information described above will be displayed on the screen for another 15 seconds. After clicking the button "OK" once again, a new round starts. On the upper right side of the screen you see the seconds counting backwards from 15 to 0. Please note that you can take more time for your decision in each round if you like. Only by clicking on the "OK" button you will arrive at the next step. However, we advise you to roughly stick to the time limit so that this phase is not unnecessarily prolonged.
Before starting your predictions, you will see 10 values of the variable X and the corresponding signals of 10 trial rounds on screen. Those values of the variable X and the signal were generated in the same way as it is done in the following 60 prediction rounds. It can be very helpful to study the information in this table carefully, especially for making your predictions in the first rounds.

[*No-Info TREATMENT:* Notes on the variable X and the signal:
The variable X can take integer values in the range of 385 to 785. The signal is an indicator for the value of the variable X in the current round. The signal thus hints at the value of the variable X in a round, but the signal contains an error that is randomly drawn in each round from a specific distribution. The distribution of the error term is the same in all rounds. The average of the error over many rounds is 0, but the error in each single round can be positive or negative. That means that the signal can be larger or smaller than the actual value of

the variable X in any round. Nevertheless, the relation between the variable X and the signal is positive. That means, on average, the higher the signal the higher the value of the variable X.]

[*Info TREATMENT:* Notes on the variable X and the signal:
The variable X can take integer values in the range of 385 to 785. In each round, the value of the variable X is randomly drawn from a normal distribution with an expected value of 585 and a standard deviation of 50. The distribution is truncated so that the randomly drawn value of the variable X can never be smaller than 385 and never be larger than 785. To illustrate the probabilities of different values, the following example shows a typical result when drawing a value from this distribution 1000 times. The height of the bars indicates how often a particular value in the range of 385 and 785 was drawn in these 1000 repetitions. This gives an impression about the probabilities that X takes on specific values in the specified ranges.

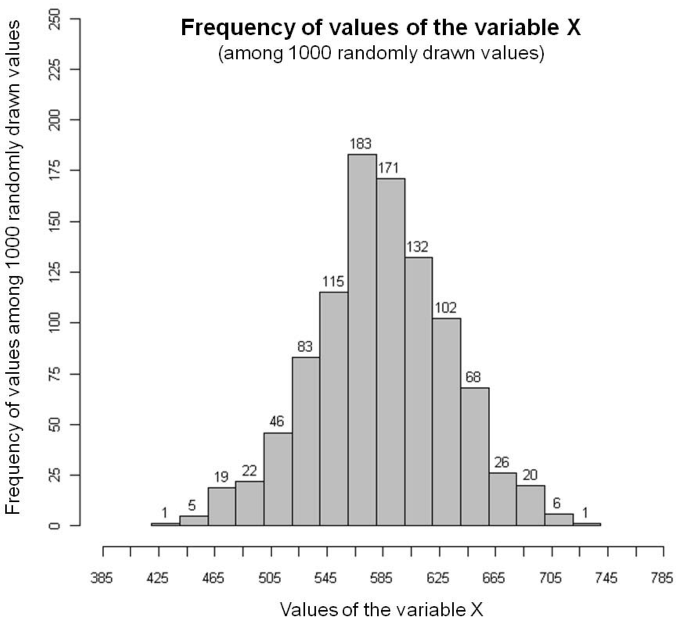

Properties of the distribution:

- Values close to the expected value of 585 are more likely and thus occur more often than more extreme values
- The farther away the values from the expected value, the smaller the probability that they are drawn
- Values smaller than 385 and larger than 785 are not possible

The signal is an indicator for the value of the variable X in the current round. The signal thus hints at the value of the variable X in a round, but the signal contains an error that is randomly drawn in each round from a specific distribution. The distribution of the error term is the same in all rounds. The average of the error over many rounds is 0, but the error in each single round can be positive or negative. That means that the signal can be larger or smaller than the actual value of the variable X in any round. Nevertheless, the relation between the variable X and the signal is positive. That means, on average, the higher the signal the higher the value of the variable X.]

Payoff for phase 3:
Your payoff in phase 3 depends on the accuracy of your prediction of the value of the variable X. Out of all 60 rounds, the computer will select one round at random at the end of the experiment (i.e., after phase 4) and you will be paid according to your prediction in this round.

If you predicted the value of the variable X in this round exactly, you earn **6.00 €**. If your prediction deviates from the actual value of the variable X, your payoff is calculated as follows:

$$\text{Payoff} = 6.00\text{ €} - 0.015\text{ €} \times |\text{Your prediction} - \text{Value of the variable } X|$$

"$|\dots|$" symbolizes the absolute value of the difference between your prediction and the actual value of the variable X. That means that it is decisive for your payoff how far your prediction deviates from the actual value of the variable X, irrespective of whether your prediction was below or above the actual value. According to the above formula, 1.5 €-Cent is subtracted from 6.00 € for each point your prediction deviates from the actual value. In general, the more your prediction deviates from the actual value, the less you earn. In case your prediction deviates only a little, the amount subtracted from 6.00 € is small; in case your prediction deviates a lot, the amount subtracted from 6.00 € is larger.

To make this even clearer, please look at the examples in the table.

Actual value of var. X	**Your Prediction**	**difference** (Prediction - Value of X)	**absolute value of the difference (deviation)**	**$0,015 \times$ deviation**	**payoff**
500	**550**	50	50	0,75	**5.25**
500	**450**	-50	50	0,75	**5.25**
650	**630**	-20	20	0,3	**5.70**
650	**670**	20	20	0,3	**5.70**
400	**500**	100	100	1,5	**4.50**
400	**300**	-100	100	1,5	**4.50**
385	**785**	400	400	6	**0**
785	**385**	-400	400	6	**0**

(As it can be seen in the last two lines of the table, you cannot make losses. In case of the largest possible deviation of your prediction from the actual value of the variable X (i.e., 400), your payoff would be 0. Such a large deviation is, however, very unlikely).

Phase 4

In phase 4, you are a trader who can buy and sell shares of an assets on a market. The market consists of 8 such traders. For phase 4, you obtain a basic amount of 10.00 €. This amount can increase or decrease, depending on your gains and losses in this phase.

In the following, all amounts are denoted in ECU. 25 ECU correspond to 1.00 Euro. At the end of this phase, your final amount of ECU will be converted to Euro and be added to, respectively subtracted from, the basic amount of 10.00 €. Phase 4 consists of 10 rounds.

Value of the asset

At the beginning of each round, each trader receives 20 shares of the asset. The value of the asset is randomly drawn in each round from a normal distribution with expected value of 120 ECU and a standard deviation of 15 ECU. However, the asset value can never be smaller than 60 and never be larger than 180. To illustrate the probabilities that the asset takes on specific values, you find a figure below that shows a typical result when drawing the asset value from this normal distribution 1000 times.

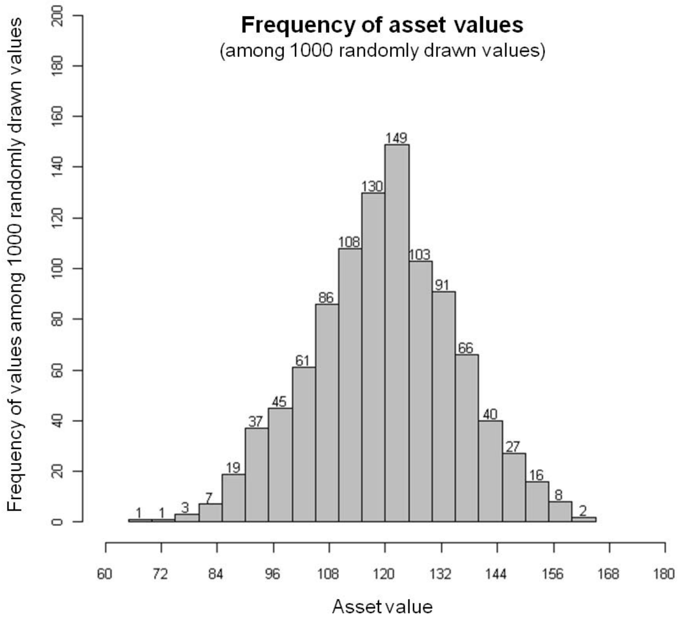

Properties of the distribution:

- Asset values close to the expected value of 120 are more likely and thus occur more often than more extreme values
- The farther away the values from the expected value, the smaller the probability that they are drawn
- Values smaller than 60 and larger than 180 are not possible

Private signal about the asset value

For each round, a new asset value is randomly drawn from the above described distribution. No trader is informed about the true asset value. However, **each trader receives a private signal about the asset value at the beginning of each round**.

The signal about the asset value is imprecise. More specifically, the **signal** consists of **the true asset value ± an error**. This error is also drawn from a normal distribution. The expected value of this normal distribution is 0, which means that on average the error is zero. In each round, the signals of all 8 traders are randomly drawn with the constraint that the average of all 8 signals correctly reflects the true asset value. However, each trader receives a private signal that can be larger or smaller than the true asset value.

To gain some impression about the relation between the signal and the asset value, you will see a list of 10 asset values and corresponding signals before the first trading round starts. The asset values and the signals were generated in the same way as it is done in the following 10 trading rounds.

Profits

By buying and selling assets you can make profits.

At the beginning of each round you and every other trade on the market receives 20 asset shares and a cash amount of 10,000 ECU on his or her starting account.

$$\begin{aligned}\textbf{Value of your starting account} = &\\ &10,000 \text{ ECU}\\ &+ [\text{20 asset shares}] \times [\text{true asset value}]\end{aligned}$$

Through buying and selling shares of the asset during a trading round, your stock of asset shares and your ECU on your account changes. If you buy shares, you will have more shares of the asset and less cash (in ECU) at the end of a trading round. If you sell shares, you will have less shares of the asset and more cash (in ECU). You can also buy and sell asset shares at the same time.

At the end of a trading round, you have a final stock of asset shares and ECU on your final account.

$$\begin{aligned}\textbf{Value of your final account} = &\\ &\text{ECU at the end of the round}\\ &+ [\text{number of asset shares at the end of the round}] \times [\text{true asset value}]\end{aligned}$$

At the end of a round, you will be informed about the true asset value and thus about the value of your starting account and of your final account. *You make profits, if the value of your final account is larger than the value of your starting account. You make losses, if the value of your final account is smaller than the value of your starting account.*

$$\begin{aligned}\textbf{Profit/Losses in one round} &= [\text{Value of your final account }]\\ &- [\text{Value of your starting account}]\end{aligned}$$

<u>Note:</u> **You make profits in one round, if you**

- **buy** assets at a price that is ***lower*** **than the true asset value**.
- **sell** assets at a price that is ***higher*** **than the true asset value**.
- buy and sell assets within one round and the **selling price is *higher* than the buying price**.

Profits or losses are not transferred to the next round. That means that in each round you start again with 20 asset shares and 10,000 ECU.

Buying and selling asset shares
Each trading rounds consists of two parts.

Part 1
In part 1, you can submit a buy offer and a sell offer. On the screen, you also see your signal for the asset value.

- A *buy offer* consists of the *maximum* buying price that you are willing to pay for one asset share and a *maximum* number of asset shares that you are willing to buy at this price.
- A *sell offer* consists of the *minimum* selling price that you are willing to accept for one asset share and the *maximum* number of asset shares that you are willing to sell at this price.

If you do not want to trade at all in part 1, you leave all text boxes empty and click "OK." If you only want to submit a buy offer, then you leave the text boxes for the sell offer empty and click "OK," and vice versa. For your decisions in part 1, you have a time frame of one minute.

As soon as all 8 traders have submitted their buy and sell offers, the trading price P for the asset shares is determined in the following way: All buy offers are aggregated to determine the total demand for asset shares and all sell offers are aggregated to determine the total supply of asset shares. The trading price P is calculated by the computer as the price at which the total demand for asset shares is equal to the total supply. By that, the maximum number of asset shares is traded at the unique trading price P.

In particular, this means for you:
You *buy* asset shares, if your *maximum* buying price is equal to or larger than the trading price P. The price that you have to pay for one asset share is the trading price P. This price may be lower than or equal to your maximum buying price. You buy at most the number of asset shares that you have specified in your buying offer. It is also possible that you buy less asset shares, if the supply of asset shares at the trading price P is not sufficient for everyone who is willing to buy asset shares at this price.
You *sell* asset shares, if your *minimum* selling price is equal to or lower than the trading price P. The price that you obtain for one asset share is the trading price P. This price may be

higher than or equal to your minimum selling price. You sell at most the number of asset shares that you have specified in your selling offer. It is also possible that you sell less asset shares, if the demand for asset shares at the trading price P is not sufficient for everyone who is willing to sell asset shares at this price.

At the end of part 1, you are informed about the trading price P, how many asset shares you have bought or sold, and what your current stock of asset shares and cash (in ECU) is. After 30 seconds (at the latest), the trading round continues with part 2.

It can happen in part 1 that no trading price P can be determined at which the total supply of asset shares equals the total demand. In such a case, there is no trade in part 1. You will be informed if this happens.

Part 2

In part 2, you can directly buy from and sell to other traders on the market.

Part 2 of each trading round consists of 120 seconds (2 minutes). Please note the extra sheet of paper in your instructions that shows a schematic representation of the trading screen in part 2. On the upper right side, you see how much time (in seconds) is left in part 2 of the current trading round. In the upper middle part, you see your signal for the asset value, your current stock of cash and asset shares, and the trading price P as well as the total number of asset shares traded in part 1.

How can you buy asset shares?

There are two possibilities:

- You submit **your own buying offer** (see point 1 on the sheet with the trading screen). To do so, you specify a **maximum price** that **you would be willing to pay for one asset share**, and click on the button "**submit your buy offer**." Your buy offer now appears in the column "standing buy offers" and is visible to all other traders. In this column, you also see the standing buy offers of all other traders. Your own buy offer appears in blue color, the buy offers of all other traders appear in black. If you want to buy more than one asset share, you can make several buy offers, one after the other.

  If your buying price that you submitted is higher than (or equal to) the minimum standing sell offer of another trader, the trade is executed immediately and you pay only this lower price. *At the time of submitting a buy offer, you just have to think about how much you would be willing to pay for one asset share at most.*

- You **accept** a **standing sell offer** (see point 2). The **best (=lowest) sell offer of another trader** is at the **top** of the list and always marked in blue. By clicking on the button "buy," you buy an asset share at this price.

How can you sell asset shares?

There are two possibilities:

- You submit **your own selling offer** (see point 3 on the sheet with the trading screen). To do so, you specify a **minimum price** that **you would be willing to accept in exchange for one asset share**, and then click on the button "**submit your sell offer**." Your sell offer now appears in the column "standing sell offers" and is visible to all other traders. In this column, you also see the standing sell offers of all other traders. Your own sell offer appears in blue color, the sell offers of all other traders appear in black. If you want to sell more than one asset share, you can make several sell offers, one after the other.

  If your selling price that you submitted is lower than (or equal to) the maximum standing buy offer of another trader, the trade is executed immediately and you receive this higher price. *At the time of submitting a sell offer, you just have to think about the price you demand at least for selling one asset share.*

- You **accept** a **standing buy offer** (see point 4). The **best (=highest) buy offer of another trader** is at the **top** of the list and always marked in blue. By clicking on the button "sell," you sell an asset share at this price.

You can simultaneously submit buy and sell offers. All buy and sell offers are for one asset share each. *Buy and sell offers that have been submitted to the market cannot be deleted afterwards.* Before phase 4 starts, you have the opportunity to practice the trading procedure of part 2 in two training rounds.

You can at most sell as many assets as you own. For buying assets you have a maximum amount of current ECU holdings available. If you submit a buy offer, your current ECU holdings are reduced by the amount of the buy offer. If you submit a sell offer, your current stock of asset shares is reduced by one.

At the end of part 2 of each trading round, you are informed about your profits (or losses) of the current round, which is calculated as described above.

Payoff in phase 4:
At the end of the experiment, one of the ten trading rounds is chosen at random. Your profits in the randomly drawn round are converted to Euro and added to the basic amount of 10.00 Euro; respectively, your losses in the randomly drawn round are converted to Euro and subtracted from the basic amount of 10.00 Euro. This final payoff for phase 4 is then added to your overall payoffs in the whole experiment.

In rare instances, it is possible that you incur losses in phase 4 that are higher than the basic amount of 10 Euro. If this is the case, you can compensate your losses in phase 4 after the experiment by working on additional tasks, so that you do not get less than 0.00 Euro for phase 4. The additional tasks consist of counting numbers. You can compensate losses of 1.00 Euro for each completed task. These additional tasks can only be used to cover your losses, but not to increase your profits.

Before phase 4 starts, you have to answer some control questions to ensure that you have read and understood the instructions correctly. Thereafter, the two training rounds for the trading procedure in part 2 start. Please try out all possibilities to buy and sell assets on your own to make yourself familiar with the trading process and the trading screen of part 2. Profits or losses in the training rounds are not relevant for your final payoff! In both training rounds, the asset value is 120 ECU and there are no signals for the asset value.

If you have any questions, please raise your arm and we will come to your cubicle to answer your questions. At the end of the two training rounds, you again have an opportunity to ask questions before phase 4 starts with the real trading rounds.

Trading Screen Asset Market