The value of information and optimal trading strategies in markets with heterogeneously informed traders[#]

Jürgen Huber[a,*] and Michael Kirchler[a]

[a] Department of Finance, University of Innsbruck

Current Version: October 2003

Abstract

With a simulation study and an experimental market we explore, how valuable information in a market is. While earlier work in this field covered this question only with two levels of information we use ten different levels to control careful for the influence of additional information. We find that additional information is mostly useless and sometimes even harmful for low and medium informed investors. The second focus of the paper is to explore the usefulness of different trading strategies. Here we find that different information levels should use differing strategies, so there is no single optimal strategy.

JEL-classification: C91; D82; D83; G14

Keywords: Value of information, Experimental Economics, Trading strategies, Information asymmetries

---

[#] The authors are grateful to Klaus Schredelseker, Matthias Bank, Matthias Sutter and participants of WEHIA 2003 and ELAB 2003 for helpful comments on earlier drafts. Financial support by the Raiffeisen-Landesbank Tirol and the Center of Experimental Economics of the Institute for Public Economics at the University of Innsbruck are gratefully acknowledged.

[*] Corresponding author: Department of Finance, University of Innsbruck, Universitätsstr. 15, A-6020 Innsbruck, Austria, Tel: +43 512 507 7554; e-mail: juergen.huber@uibk.ac.at.

1 Introduction

Will an agent in a market be able to improve his expected return by gathering more and better information? Instinctively we are inclined to say “Yes, sure”, but in this article we will see, that the answer is not that simple.
The marginal benefit of additional information is usually assumed to be positive. While this assumption is true in a single-person context, as pointed out by Blackwell (1951), the situation may be more complicated in a multi-person context. In particular with respect to financial markets, it is widely believed, that traders with more information make better decisions and therefore gain higher profits. However, game theory reveals that “having more information … can make a player worse off” (Gibbons, 1992, p. 63). Even though the game-theoretical properties of markets are widely realised, little attention has been paid to potential consequences of these properties.

Apart from the efficient market hypothesis in its strongest form the existing literature, models and experiments covering the value of information in markets mostly conclude, that additional information will make the possessor of information better off. However, all of these studies have a major shortcoming: They compare only two levels of information – uninformed vs. informed.[1] Their common result, that the informed can outperform the uninformed traders, is no surprise. Until now, little consideration has been given to the impact of heterogeneously informed traders on the relation of information and return in capital markets. Yet it is a fundamental characteristic of modern stock markets that different agents receive different information signals.

The purpose of this paper is to extend the current research on the importance of information in markets by introducing more than two levels of information. We think that the design of our markets is simple enough to permit the computation of an equilibrium return distributions, yet rich enough to capture the forces at work in financial markets.
The rest of the paper is organized as follows: In Chapter 2 we provide an overview over related literature and ideas. In Chapter 3 the research question will be clarified before

[1] One remarkable exception is a paper by Diamond and Verrecchia (1981) that will be discussed in Chapter 2.

turning to the design of the market. Next we will examine the used trading strategies and then present the results of the simulation study. The analysis will be continued by changing trading strategies in the simulation to derive an equilibrium. In Chapter 4 the experimental results are presented before the paper is concluded with Chapter 5.

2 Related Literature and Ideas

Since decades the efficient market hypothesis (EMH, see Fama (1970)) is one of the cornerstones of finance. If it holds, prices “fully reflect all available information at all times” (Fama 1970, p. 385) and gathering information is useless, as all information is already incorporated in the market prices. Grossman (1976) attacked this position by formulating his information paradox: If prices fully reflect all available information, nobody would gather (costly) information. But if nobody gathers information, how can prices possibly reflect all (or any) information?

Grossman and Stiglitz (1980) solve this puzzle by assuming an efficient level of inefficiency (noise) in the market, which allows some traders to increase their return by gathering information. However, with only two types of traders – informed and uninformed, their model is rather about asymmetric, than diverse information. In equilibrium the extra return of these agents is assumed to be exactly their information costs. As a result the net return after information costs is again the same for all traders. If we realistically assume positive information costs, this net return would have to be below the market return, as any information costs would lower the average net return. The higher the information costs and the more traders gather information, the lower their expected net return will be. We agree with Malkiel (2003a, p.2), who argues, that “clearly all stocks have to be held by someone and if certain investors achieve above-average returns, then it must be the case that other investors are achieving below average performance … after accounting for the additional expenses of active management, most investors must underperform the market average.”

This leads to the starting point of our research: If the random walk hypothesis holds, a trader who does not gather any information but trades randomly in time can expect to earn the market return. There is no reason to assume any systematic over- or underperformance, if she really chooses her shares randomly, e.g. by throwing a dart arrow at a quotations list. This would make a random trader the best net performer in a market, as she would receive the market return, while all traders gathering information would have an expected net return below the market return.[2]

This could explain why high-paid funds managers are regularly not able to beat a broad market index, as shown in studies across the globe. On average about 70 percent of actively managed stock market funds were outperformed by the market over a ten-year period, for bonds the number is even higher at 90 percent. Over the past ten years the median actively managed fund has produced annual returns 175 basis points lower than the index (Malkiel 2003a, p. 4 and 9). In less scientific competitions of funds managers against cats, apes or school children, the "outsiders" are regularly able to outperform the professional managers, as the former basically act randomly, while the professionals process information.

We tried to find a solution to this puzzling result in the existing literature, but most of the papers we found were of little help. We found a number of papers, some of them about experimental studies, exploring the relationship of information and return in a market, but they all distinguished only two information levels (informed vs. uninformed). Their common result, that the informed can outperform the uninformed is no surprise (see for example Ackert et al. 2002, Sunder 1992, Copeland/Friedman 1992, Haltiwanger/Waldman 1985).

Diamond and Verrecchia (1981) develop a market model with a large number of heterogeneously informed traders, but their information system is independent, as it includes many diverse sources of information, and their focus it the markets' ability to aggregate information in this "noisy" environment. For our analysis this model is of little help, as the precision of information is identical across traders, since each trader has the same prior beliefs and is endowed with private information of the same precision. The focus of our research question is how valuable information of different extent and precision is to the owner of the information.

No accepted analytical or experimental study that we are aware of examines market dynamics when there are more than two levels of information. Reality, however, most definitely is characterized by a multitude of different information levels of individuals.

[2] As Hishleifer (1971, p. 573) points out, there is an incentive for individuals to expend resources in a socially wasteful ways in the generation of information.

Here we see a major shortcoming in present studies and the purpose of this paper is to extend this line of research by introducing more than two levels of information.

One author covering the relationship of information and return in a market with more than two levels of information is Schredelseker. He approaches the problem analytically (Schredelseker 1984) and with a simulation study (Schredelseker 2001). In his papers Schredelseker uses a different line of argumentation, than we did above, as he does not take information costs into account, but his main conclusion is the same. His starting point is the above-average return of insiders: If we have a less than strong form efficient market, insiders will be able to gain above average returns. As we saw above, a random trader can expect the average market return. This leads to another puzzle: If some traders gain more than the market return and the uninformed receive the market return, who is below the average? The only possible group are the average informed traders. The resulting relationship between information level and return looks as follows:

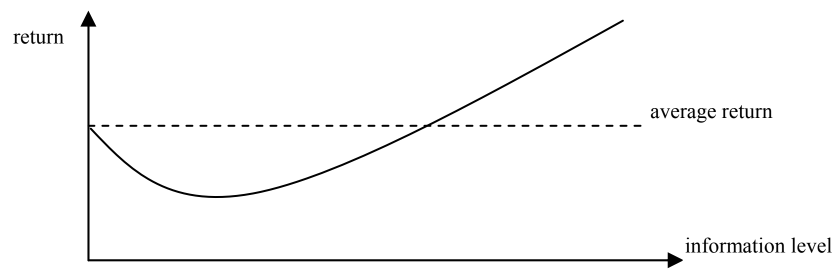

Fig. 1. Rate of return per information level as assumed by Schredelseker (1984, p.51)

The intuition behind this stunning result is quite simple: Whenever an average informed agent trades in the market he takes a bet against a better informed person. As a consequence he will loose on average. In a market context this means being below the average return. That relationship is well established in game theory: If I play against a person who is able to foresee my moves better than I can predict hers, I should make myself unpredictable by playing randomly. The same is true for markets, and making himself unpredictable is exactly, what an uninformed trader does: He chooses his stocks randomly. Consequently there is nothing foreseeable and therefore nothing exploitable in his trading strategy. Again, the below-average performance of professional investment funds managers, which are in most cases not able to beat a broad market index, suggest that gathering information could, in fact, be futile to a large extent (Malkiel, 2003a and 2003b).

Strong and Walker (1987) also examine the relationship between information level and return, but their argumentation is not as explicit as Schredelsekers. Their result, however, is quite close to his, as they state, after finding that the low informed investors systematically loose to better informed agents, that low informed traders have two ways to avoid these losses: "*They can refuse to trade with the more informed individual and/or insulate themselves from the trading activities of the more informed by adopting a passive 'buy and hold' strategy*." (Strong/Walker 1987, p.91).[3] They also argue that by refusing to trade on information (passive strategy) a trader can improve his situation. In Fig. 1 this can be interpreted as shifting from an average information level (earning less that the average return) to a strategy using no information (with an expected return equal to the market return).

3 Research question

With this study we want so explore two questions: First we want to find out, how valuable information of different precision is for agents in a market. Second we analyse, which information processing strategy is best for each trader and whether there are differences for different information levels. .

The model developed by Schredelseker (2001) in his simulation study seems very well suited to investigate the return of heterogeneously informed agents in a market with more than two different information levels. We therefore chose his model as starting point for our study and adapted it for our research questions. To gain reliable results we chose a two-pronged research design, testing the same model experimentally and in a simulated market. The results will be presented separately, but first the design of the market will be illustrated.

---

[3] The "no-trade" strategy can be countered by comparing the return in risk-free and risky assets. If the risk-free rate is for example 3 percent and the average return in the stock market is 8 percent, traders have an incentive to stay in the market, as long as their under-performance is less than 5 percent, as their net return will still be better than the alternative.

3.1 Design of the market

To investigate the usefulness of information in a market we chose the simplest market design that is still capable to capture the main characteristics we need. Our economy has only one asset that can be traded by going long or short in each round. The rounds are independent, and each round consists of four steps: First the traders are endowed with information about the intrinsic value of the asset. Then the agents post their reservation price (henceforth "bid") for the asset. In a third step the ten bids are sorted from the lowest to the highest. To clear the market the median becomes the market price and all traders with a lower bid are sellers (short), while all traders with a higher bid are buyers (long). Traders having bid the median are neutral in this round.[4] To ensure the zero-sum property of the market and to make sure, that the net supply of papers is zero, scale selling is used if the number of buyers and sellers is not equal. With this design we have a double-auction market without spreads, where all traders act as market makers. The market price thereby reflects supply and demand and is an endogenous variable. In the fourth step of each round the individual payoff is calculated by comparing the market price with the intrinsic value according to the formula

$$R_{j,k} = \frac{Bid_{j,k} - P_k}{\left|Bid_{j,k} - P_k\right|} \cdot \left[V_k - P_k\right].$$

$R_{j,k}$ stands for the return of trader j in round k, $Bid_{j,k}$ is the posted bid of trader j in round k, P_k is the market price and V_k the intrinsic value in round k. A buyer makes a profit if the intrinsic value is higher than the market price. If the intrinsic value is below the market price, he makes a loss. A seller receives a positive payment, if the intrinsic value is below the market price, and vice versa. For example, if the market price is 5 and the intrinsic value is 6, each buyer gains 1, while each seller looses 1.

The intrinsic value and the information system:

In each round the intrinsic value of the asset is given by the sum of ten Laplace-coins showing either 1 or 0 with the same probability. The coins represent different brackets of the total information. The distribution of the intrinsic value is thus generated by a binomial process of ten steps leading to a binomial distribution with a mean of five and

---

[4] If the bids are, for example, 0-3-4-4-**5-6**-7-7-7-8, a market price of 5.5 prevails. If the fifth bid were 6 instead of 5, the market price would have been 6 and the two traders bidding 6 would be neutral.

a standard deviation of 1.58. This generation process for the intrinsic value is known to the traders.

The ten coins are used to create a heterogeneous information structure among the agents. Trader Ix knows x of the ten coins, with x ∈ {0, 1,…, 9}. I0 therefore knows none of the coins, I1 knows the realization of the first coin, I2 knows the first and the second coin, etc. until I9 who knows nine of the ten coins. The information system of our market is cumulative – a better informed trader always knows all the coins a worse informed trader knows plus some additional information. If the coins are understood as representing different brackets of all the factors relevant for a share price (e.g. the first coins representing inflation outlook, economic growth, industry sales outlook, etc.; and the last few coins representing secret product developments or the retirement of a key executive), it seems rational and realistic, that the low informed traders tend to have similar information (from newspapers, newsletters and TV), represented by a cumulative information system, than all having strongly diverging information represented by an independent information system. This makes the last few coins very exclusive information that can be considered insider information.

We decided to use a design with ten traders, as the concept of (beta) risk deciles is well established in finance. Analogously we introduce information deciles represented by the ten different information levels.[5]

3.2 Trading strategies:

In the simulation only two different trading strategies were used to keep the analysis as simple as possible. The most straightforward strategy is for each trader to use the information he has to estimate the intrinsic value of the asset. This strategy, which we call “active information processing strategy”, or “information processing strategy”, is equivalent to the fundamental analysis used by many participants in real stock markets. The trader sums up the value of all the coins he knows and adds the expected value of all other coins. If the six coins I6 sees show 110101, he knows that the intrinsic value is between four (if the four coins he does not see all show 0) and eight (if they all show 1).

[5] In a real market the low information classes I0 to I4 would probably be the millions of small investors who do not have access to special information, while the highest information classes I8 and I9 may be a small number of insiders with billions to trade with. Professional funds managers would probably found in the categories I5 to I7.

The expected value of the four unknown coins is two, so his best estimate of the intrinsic value is six. This can be generalized for all information levels:

$$Bid_{j,k} = \sum_{i=1}^{j} c_{i,k} + [(n-j) \cdot 0.5],$$

with $Bid_{j,k}$ representing the bid for the trader j in round k, $c_{j,k}$ the coins known to trader j in round k, and n the maximum number of coins (10).

The second possible trading strategy in our model is to act randomly ("random strategy"). We decided to use this strategy instead of a passive investment strategy, as it is more flexible and sometimes easier to employ in a real market. Here several approaches would be possible – for example choosing each value form 0 to 10 with a probability of 1/11, or choosing only 0 or 10 with 50 percent probability. We decided to use the simplest strategy available: A random player posts a bid of zero or ten, each with p=0.5.

We want to stress several design features, before we come to the results: (i) in the simulation as well as in the experiment the information level of each agent was kept constant during the whole session. (ii) We have absolutely no information or transaction costs. (iii) Each agent was required to post a bid each round. It was therefore not possible to abstain from trading,[6] and (iv) we have a pure zero-sum game where zero is the benchmark (the market return) for each trader. All profits/losses reflect only the result from market trading. If information costs were included, the return of better informed traders would drop depending on the extent of information costs, while uninformed traders would still pay nothing for information.

4 Simulation Results

We will now analyse the simulation in depth and turn to the experimental results in Chapter 5. For the results of the simulation study presented below the average returns of

[6] We are aware of the no-trade theorem, but as pointed out in footnote 3, the higher return in a stock market can induce traders to stay in the market and continue trading even when their return is below the average. They will stay in the market if the return there is still higher than any other alternative.

all 2^{10}=1024 possible sets of coins have been calculated. All of the results presented show the average of twenty simulation runs we did for each result to reduce the influence of chance in the return of random traders.

As the sum of the ten coins gives the intrinsic value, and as a coin is never shown wrong, everybody will agree, that the more coins a trader knows, the higher is his information level. It is obvious, that a person knowing five of ten coins is better informed, that a person knowing none. The question we want to answer is, whether the better informed can expect a higher return in the market than the uninformed.

In Fig. 2 we see the resulting expected returns for the ten traders when all but I0 process their information actively. The assumption of a positive additional value of information holds for the fourth and all additional coins, but it is obvious, that the first three coins can not improve the return of the trader, but worsen it. I6 is the first who is able to outperform the uninformed I0. Not the worst informed trader (I0), but the average informed I3 has the lowest return in the market.

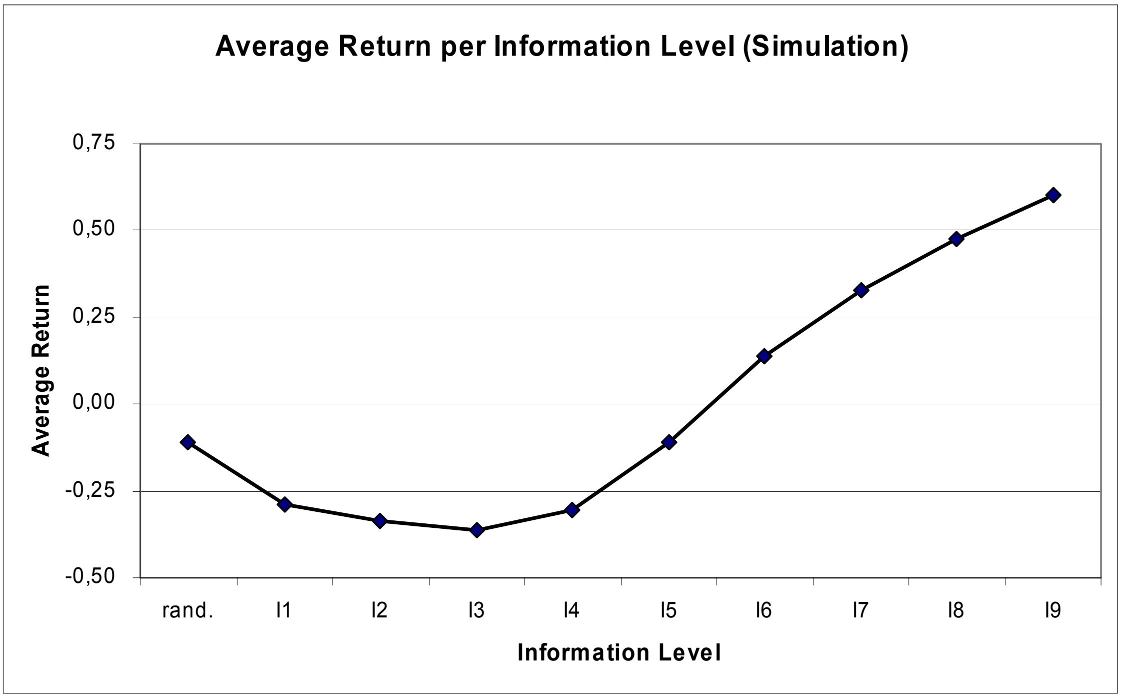

Fig. 2. Rate of return per information level in the simulation approach

Confronted with this result we see two questions to answer immediately: Why I0 is below the average of zero, and more importantly, why and how additional information can worsen the return of I1, I2 and I3.

We argued that in a market a random trader can expect the market return, as there is no reason to expect systematic errors in any direction. This is only true in a liquid and broad market, where every single trader is a price taker. In our narrow market this condition does not hold, as every trader represents one tenth of the total market and therefore has significant influence on the market price. Whenever I0's bid is zero, he is a sure seller and thereby possibly lowers the market price and vice versa when his bid is ten, making him a sure buyer. This influence on the price is mostly to his disadvantage and decreases his return to the slightly negative average number we see in Fig. 2. In a market with thousands of traders his return would be (very close to) zero.

Now to the more interesting question of the negative value of additional information for low to average information levels: The key to solving this puzzle is biased information. If information is unbiased, a part of the information shows basically the same picture as the total information. If the ten coins show for example 0101011010, an uninformed trader estimates a value of five, a trader knowing the first four coins (0101) also estimates five, and the same holds for a participant knowing the first eight coins if the traders use active information processing. In cases of unbiased information, as above, most traders estimate the same price, which will also be the market price. At this price nobody will loose or gain and the information is basically useless. While most sets of coins are of this type, some are cases of biased information.

Here a part of the information shows a different picture than the total information. Let us consider a case where the coins show 0000011111. As in the example above the intrinsic value is five, but some traders may be mislead by the information they get. As this is one of the core points of our study we will take a close look at every single trader and his bid with active information processing. To make the analysis as straightforward as possible we apply a passive strategy for I0 instead of a random strategy.[7] The only information I0 has is the number of coins. He therefore estimates the intrinsic value as 0.5x10 = 5. I1 knows the first coin and her bid is therefore 0.5 lower at 4.5. Subsequent estimates of the value decrease until I5 who has the lowest estimate of 2.5. I6 to I9 see additional coins showing '1', so their bids are higher.

---

[7] The results would still be the same with a random strategy, but we would have to distinguish two cases, which would make it probably harder to follow the analysis.

Table 1

Market Analysis with biased information (coins show 0000011111)

	I0	I1	I2	I3	I4	I5	I6	I7	I8	I9
bid	5	4.5	4	3.5	3	2.5	3	3.5	4	4.5
market side*	B	B	B	S	S	S	S	S	B	B
profit/loss	1.25	1.25	1.25	-1.25	-1.25	-1.25	-1.25	-1.25	1.25	1.25

* B=buyer/long, S=seller/short

The market clearing price with these bids is 3.75. Five traders (I0, I1, I2, I8 and I9) have posted bids higher than this price and are thus buyers, while the other five traders sell the asset.

Profits and losses are calculated by comparing the price (3.75) with the intrinsic value (5). In this case the asset was sold too cheap, as it is worth more than what it was traded for. Therefore each buyer receives a profit of 1.25, while each seller looses this amount. The reader will have noted, that here the average informed traders loose, while the low informed and the well informed make a profit. The average informed traders all process the same information (the first few coins which all show 0) and are thereby systematically mislead, lowering the price to their own disadvantage. We could speak of an "information risk" that traders take on, when they gather information, as this information could be biased. The information risk is zero for uninformed traders, increases with additional information until a certain extent of information is reached, where the information risk becomes lower again, as the information known is a very large part of the total information, so it has to be representative and can not be strongly biased.

We can conclude, that an uninformed trader knows too little to be systematically mislead, while the well informed know enough to be on the right side of the market. Only the average informed traders who all use the same information make the same estimation errors and influence the market price to their disadvantage. For them information is either useless (when it is unbiased) or harmful (when it is biased). Even though cases of biased information are rather rare, they contribute greatly to the final outcome, as the profits and losses in these cases are quite large.

4.1 Changing trading strategies

Traders are neither stupid nor do they like to loose money systematically. If a random player looses just -0.09, while for example I3 looses -0.37 by using the information content of her three coins, she will at some point just stop using her information and switch to a random strategy. The result of such a change in her information processing strategy can be seen in Fig.3: By ignoring her (correct) information, the trader can significantly improve her performance in the market. She is even able to outperform the better informed I4 and I5. By ignoring her information, I3 becomes a second random trader, as signified by "rand." in Fig. 3.

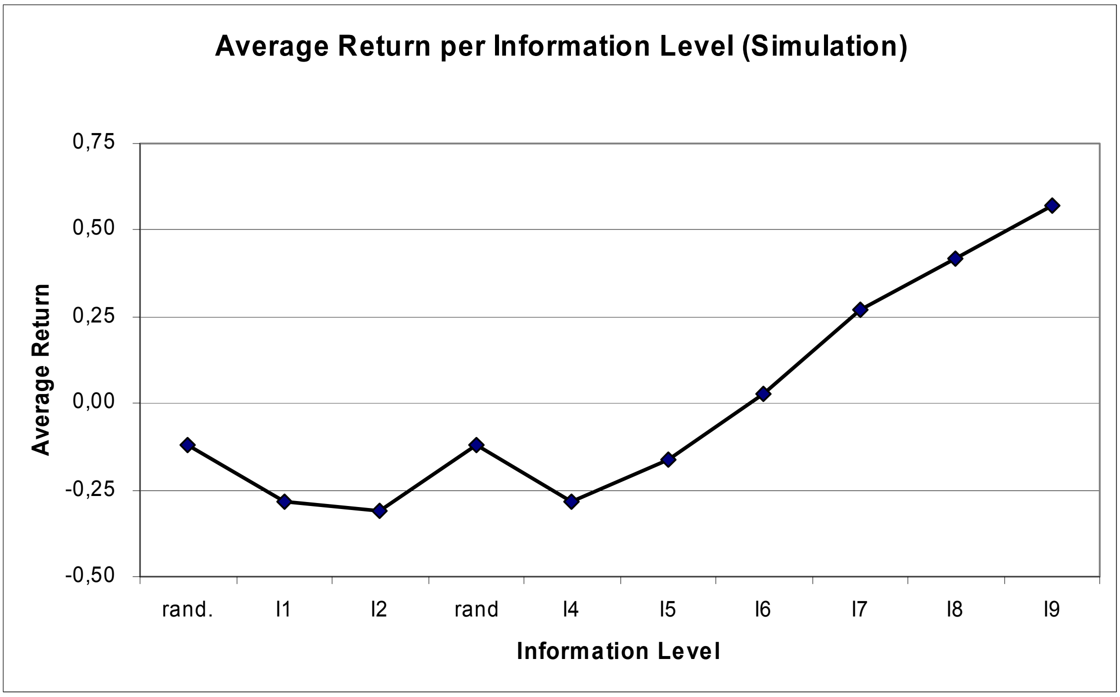

Fig. 3. Rate of return per information level in the simulation with I3 using a random strategy

The rationale behind this result has already been outlined above: I3 surely knows more than a random trader, but when trading in the market she is always betting against still better informed traders, who can systematically exploit her mistakes.

Her limited information (equivalent probably to the information provided by newspapers, stock market TV and newsletters) does not show the whole picture and may therefore be biased. As we have seen above especially the average informed traders fall into this trap, while a random player is not exploitable. Sometimes he will be right, sometimes wrong, but he does not make systematic mistakes. This result is in line with

Malkiel (2003a, p. 10), who concludes that “Investors are likely to achieve far higher returns by employing a passive indexing strategy than they are likely to achieve from active portfolio management.” A passive indexing strategy is just another way not to process any information and leads to similar results as a random strategy.

In time more and more traders will realize, that their return is below the market average and some will switch from information processing to a random strategy. However, we found, that trading strategies have a decreasing marginal return. When more and more traders start acting randomly the noise in the market increases and the better informed traders are able to profit from this – at the cost of the random traders. Fig. 4 shows, what happens if all traders who lost money in the basic simulation (I0 to I5) trade randomly: They all loose approximately the same amount of money, but with an average of -0.40 the loss is much larger than it was for the single random trader in the basic scenario. The losses of the random traders are now even higher, than the highest loss was in the basic scenario (-0.37 for I3). Please also note the changed scale, as the profits for I9 increased from 0.63 in the basic scenario to 0.83 in the case of six random traders.

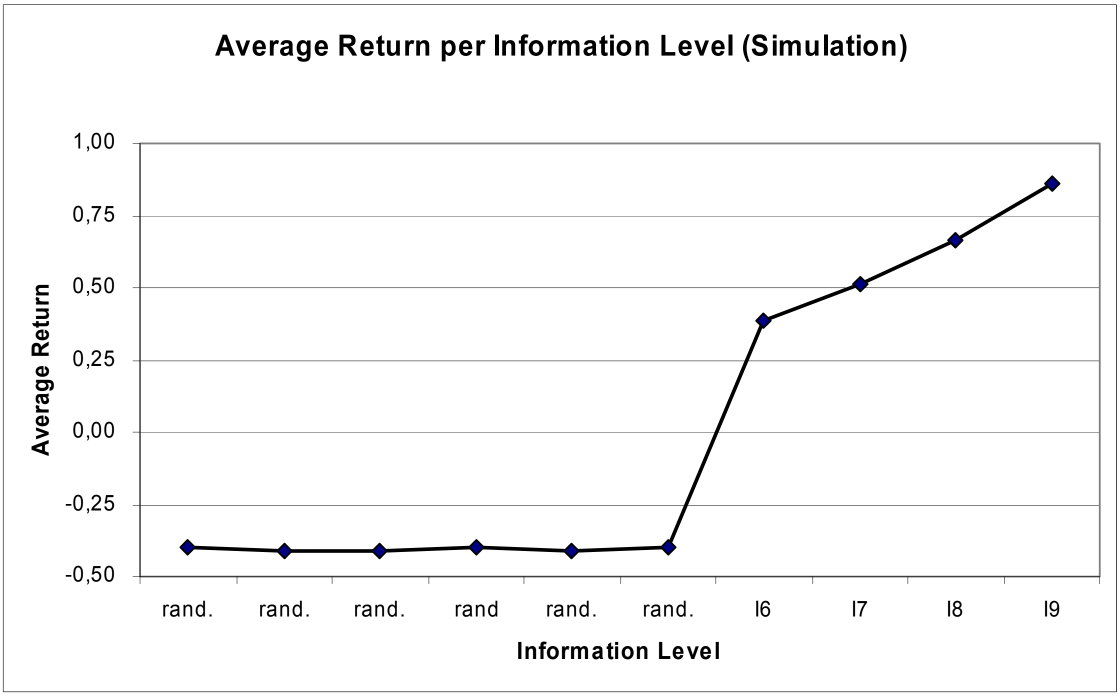

Fig. 4. Rate of return per information level in the simulation with I0 to I5 using a random strategy

The negative marginal return of the random strategy is clearly visible. The net loss increases from -0.09 when one participant trades randomly to -0.10 for two traders and a staggering average loss of -0.40 with six traders acting randomly. When the number of uninformed traders increases, the price becomes noisier and the potential return from being informed increases. To benefit from a random strategy it is important, that not too many traders use this strategy.

Let us take this comparative-statically analysis one step further: Some of the average informed traders may realise, that they can do better than shown in Fig. 4. If, for example, I3 decides to start processing her information again, she can improve her performance from -0.37 to even a small profit of 0.03. This is shown in Fig. 5. It is also remarkable, that the return of the five remaining random traders improves from -0.40 to -0.30, while the best informed loose some of their profit. This is due to decreasing degree of noise in the market when I3 switches back to processing her information.

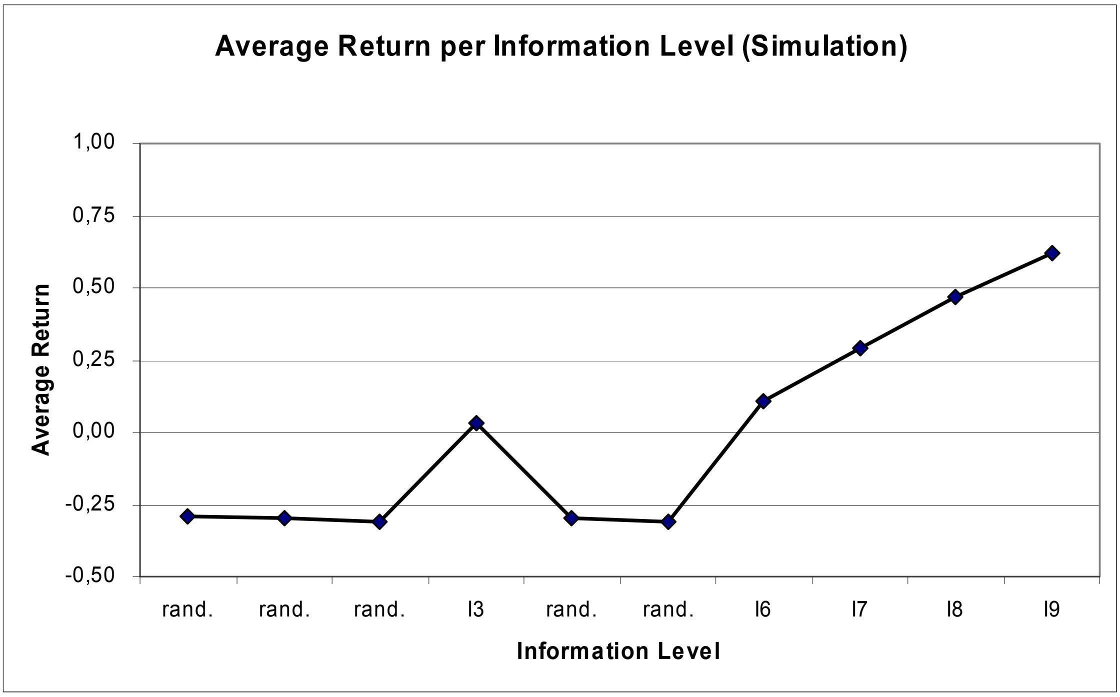

Fig. 5. Rate of return per information level in the simulation with I3 switching back to information processing

While processing information was damaging for I3, when all traders did it, it is beneficial to her, in case that many traders act randomly. This shows, how dependent the optimal strategy and the result of one trader is on what other participants in a market are doing, as is typical for games. A rough rule for low to average informed agents that we could derive from this analysis would be to "use a different strategy than all the others do." If most traders process information, it is beneficial to use a random strategy, but if many traders use a random strategy, it is better to process information even if a trader has only a rather low information level. For insiders it is a dominant strategy to use their information.

The logical question is, whether an equilibrium exists, where no trader can improve his situation given the strategies of all other traders. It is clear, that the situation in Fig. 5 is not an equilibrium, as traders like I4 or I5 can also hope to improve their return by switching back to an active information processing strategy.

Above we mentioned the decreasing marginal return for the random strategy. It should be noted, that the same can be found for the active information processing strategy: If only one agent (the best informed) uses information, while the nine others trade randomly, his profit is 1.20. The average profit drops to 1.12 when two agents process their information actively, 0.93 for three agents and a modest 0.09 if six agents process their information actively.[8] As we have a zero-sum-game the average profit is naturally zero, if all traders process information actively. Table 2 shows the complete set of average and marginal returns for both strategies depending on the number of traders using the strategy. Remember that the sum of traders using the active or the random strategy is always ten, as we have ten traders in the market. So if we look at four traders processing their information actively and each of them earns 0.60 on average, we see that the six other traders (acting randomly) each loose -0.40.

Table 2

Average and marginal return of trading strategies with respect to the number of traders using the strategy

traders using stategy	1	2	3	4	5	6	7	8	9	10
average return (active)	1.20	1.12	0.93	0.60	0.22	0.09	0.05	0.03	0.01	0.00
marginal return (active)	1.20	1.06	0.78	0.38	0.02	-0.24	-0.27	-0.30	-0.31	0.00
average return (random)	-0.09	-0.10	-0.12	-0.14	-0.22	-0.40	-0.40	-0.28	-0.13	0.00

[8] For this analysis we started with the best informed agent I9 and subsequently added the next best informed agent step by step.

For the random strategy marginal and average returns are always the same so we display only one line for them. For the active information processing strategy the marginal return is lower that the average, as the last (marginal) trader processing information is always worse informed than the ones that already used the strategy.
Let us examine an example to make the results shown in the table clearer: If three traders use a random strategy each of them faces an expected average loss of -0.12. The seven other traders, who process their information actively, make an average profit of 0.05. But not all of them win. The worst informed active trader makes a loss of -0.27 (shown in the marginal return of the seventh active traders). If this agent decides to switch to a random strategy (becoming the fourth random trader) he and the other random traders now have an expected loss of -0.14, while the average return of the remaining six traders processing information actively increases to 0.09.

We see that the marginal return with active information processing is regularly much lower than the average, as worse informed traders start to use their information. For the sixth trader processing information actively the marginal return even turns negative and it falls further with more traders trading on the information they have.

To derive an optimal strategy for herself a participant has to compare the marginal returns of each strategy. For the random strategy marginal and average returns are the same, while the marginal return with active information processing is regularly lower than the average. It is obvious that not all traders will act randomly, as the first active trader will be able to earn a profit of 1.20. Several other agents will use their information until we reach the sixth agent. When trading randomly he would loose -0.22 (as the fifth random trader) each round, if he processes his information the loss would increase slightly to -0.24 (as the sixth active information processor).
It should be noted that the marginal returns of both strategies are nearly the same in this equilibrium with -0.22 for the random strategy and -0.24 from the information processing strategy. Both strategies (and many others that we can not show here) have their place in the market.

We therefore find an equilibrium in our setting with only two allowed strategies with the first five participants trading randomly, while the five best informed traders process their information actively.

The resulting returns are shown in Fig. 6. This result can be considered a rational expectation equilibrium as no agent can improve his situation by changing his strategy and prices are market clearing for each possible realisation of the coins.

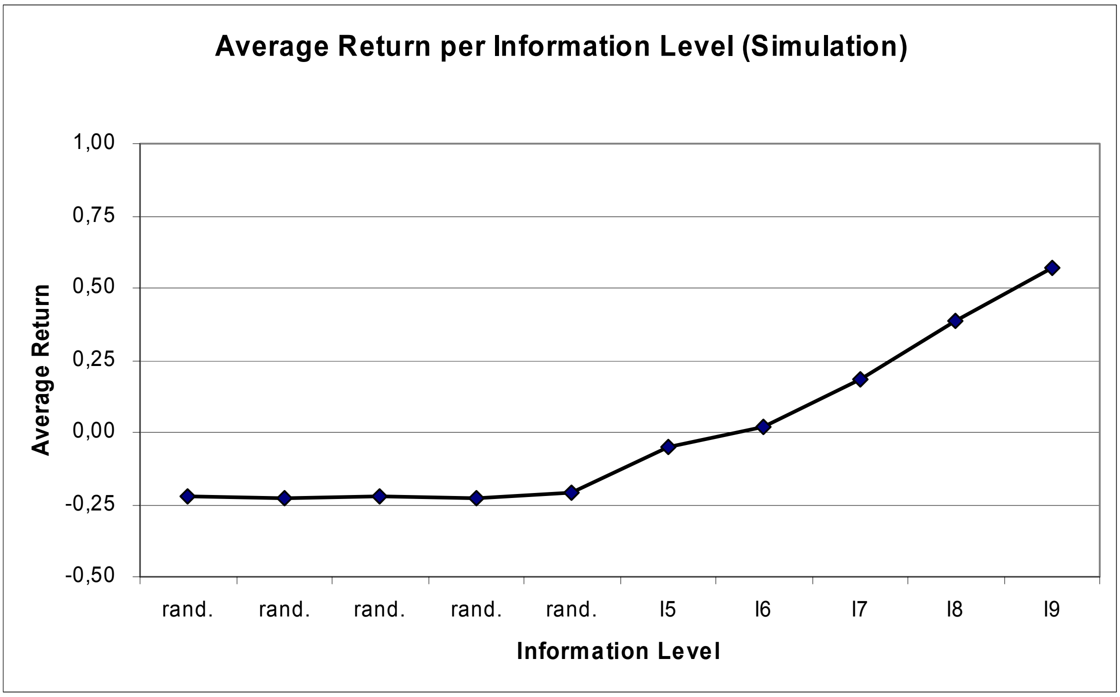

Fig. 6. Rate of return per information level in the simulation approach with equilibrium strategies

If we would allow more strategies (passive, contrarian, technical trading strategies), the task of deriving the equilibrium would become more complicated, but we think that there always exists at least one in this sort of closed market.

The decreasing marginal return of active as well as random trading strategies leads to the conclusion, that in a market with numerous traders we have an equilibrium with some traders processing information, while others trade randomly, and again others may follow any other rule that you can think of. In equilibrium the marginal returns for all these strategies should be the same and slightly below the market average. Only for the really huge investors with several billion euros to invest will it be profitable to spend money to gather information, which should allow them to gain excess returns.[9] Even

[9] In real markets several factors make the analysis probably too complex to test it empirically: Traders do not necessarily stay in the same information level, they may improve over time. The permanent entry and exit of traders in the market increases the dynamic further. In addition a correct analysis would

though we chose a very different methodological approach than Arthur et al. (1997), we can confirm their core result that several different trading strategies (even technical trading rules) may be viable in a market.
To gain some insights into the behavior of humans in this market environment and to test the reliability of the simulation results we decided to run an experimental study with the same design as the simulation.

5 Experimental Results

While the simulation was run with all 1024 possible sets of coins this was not feasible in the experiment due to time constraints. We therefore chose twenty realisations that were representative for the whole sample with respect to mean, standard deviation and expected return with active information processing.[10] Trading in the experiment was therefore implemented in twenty independent rounds, with the same set of coins in each of the experimental sessions.

The experiment was conducted with 63 students of the University of Innsbruck in seven sessions. A computer-generated random trader I0, choosing zero or ten with a probability of 0.5 each, was added to each group of nine students (I1 to I9). This approach was chosen, because we did not want to frustrate a participant by knowing nothing for the twenty trading periods.
As it is common in experimental economics we used real cash as incentive for the participants. Each participant was endowed with a certain amount of starting capital, which depended on his information level. The final payout was derived by adding/subtracting his profits/losses from trading. Each session of the experiment took about 75 minutes and the average payout to students was € 14. The computer labs at the University of Innsbruck are equipped with sliding walls to ensure that participants cannot communicate with each other during a session. The experiment was

require a separate analysis for every single asset, as some traders being extremely good informed about one asset (e.g. Bill Gates about Microsoft) may know little or nothing about other assets.

[10] To test this we ran a simulation with nine traders processing their information actively and I0 trading randomly. The resulting expected returns were very close to the ones found in the complete simulation and the correlation coefficient between the two sets of returns was 0.99.

implemented with the software z-Tree[11] which was developed especially to conduct experiments in economics.

At the start of the experiment each trader is randomly assigned a different information level which is fixed for the whole experiment. This design is used to observe possible learning effects.

Each trader knows his own endowment and the general distribution of information levels. The structure of the information system was also explained to the participants before the start of the experiment. Overconfidence, regularly observed in many real markets, should therefore be a minor problem in our experiment, as participants know, how well they are informed relative to others.

The experimental results presented below confirm the conclusions we drew from the simulation. In fact, the average returns derived from the experiment shown in Fig. 7 are very close to the equilibrium returns of the simulation we saw in Fig. 6. Each point in Fig. 7 represents the average return of a trader in one of the seven experimental sessions, while the line shows the average across all sessions.

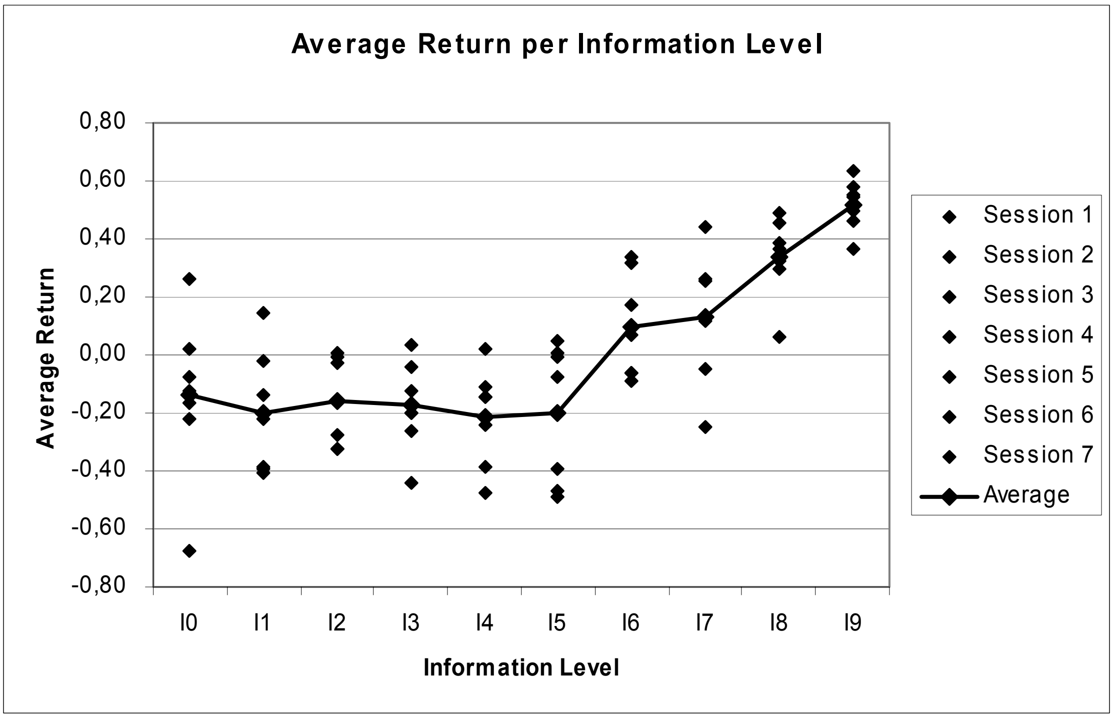

Fig. 7. Rate of return per information level in the experiment

[11] © Urs Fischbacher (1999), Zurich. We are very grateful to Urs Fischbacher and Matthias Sutter (University of Innsbruck) for helping us with the programming. We could not have done it without their help.

This confirmation of the simulation results was very encouraging for us. From the experiment we even got some insights that we first missed in the simulation. For example we see that the random strategy (I0) is probably the riskiest one, if we measure risk by standard deviation.[12] While one computer generated randomly trading I0 with luck won on average 0.28 per round (only some traders knowing six or more coins were able to do better than that), the worst trader with a net loss of -0.68 per round was also an I0. The return distributions of I8 and I9 have the lowest variances, which leads to the conclusion, that being among the best informed traders not only offers the highest returns, but these returns are also very stable and therefore related with lower risk.

What surprised us most was the resemblance of our experimental results with the equilibrium solution of the simulation presented in Fig. 6. Does this imply that our participants traded optimally? Some may, but the analysis of the experimental data suggests, that others were not able to grasp the rationale of the experiment. With only twenty rounds of trading chance also played a role, as exemplified by the above mentioned variance of the return of the random trader I0. To understand the trading behaviour we calculated the correlation coefficients between the actual bid posted by the traders and the expected bid with perfectly active information processing. The results are depicted in table 3.

Table 3

Correlation coefficient between actual and expected bid for each information level

Information level	I1	I2	I3	I4	I5	I6	I7	I8	I9
Correlation	-0.03	0.38	0.40	0.40	0.58	0.81	0.78	0.74	0.85

Generally the correlation is increasing with the information level. The well informed traders used their information when forming their expectations, while less informed traders often ignored the limited amount if information provided to them. Some traders really followed a random strategy (e.g. I1), while most average informed traders used their information to some extent. This result is consistent with the equilibrium solution of our simulation, but that does not necessarily imply, that the participants of the experiment were able to approach the equilibrium because they understood it – the

[12] The standard deviation for I0 is 0.29, while it is only 0.14 and 0.09 for I8 and I9 respectively.

trading data shows, that some average informed really switched to a random strategy after loosing for some rounds, while others were obviously not able to find a better strategy, even if they recognized their systematic losses. Still it is remarkable, that individual mistakes, which surely existed, cancelled each other out even in this very narrow market, leading to the strong result we found.

To explore the extent of learning effects deeper the performance of the traders was compared in the first vs. the second half of the experiment. Table 4 shows the results of this analysis for the first five information levels, which all lost on average in the first half. Four of them were able to improve their performance in the second half of the experiment, with the change being statistically significant for I4, while only I5 had a slightly worse return.[13]

Table 4

Learning effects in the experiment

Information level	I1	I2	I3	I4	I5
round 1-10	-0.22	-0.19	-0.22	-0.47	-0.19
round 11-20	-0.18	-0.13	-0.13	0.05	-0.21
Change	0.04	0.06	0.09	0.52*	-0.02

* statistically significant at a level of 1% (Wilcoxon Signed Ranks Test, N=7)

The improvement in the average profits of traders with information levels from I1 to I4 is due to a change from a mainly active information processing strategy in the first half of the experiment to a random strategy in the later rounds. With some caution this change in strategy can be interpreted as "learning": Some traders realised, that their low level of information did not suffice to gain above-average returns in the market and consequently they ignored it.

Another indication of learning in the experiment can be found when looking at "market efficiency". The mean average deviation of the market price from the intrinsic value can be considered as a measure for market efficiency – in an efficient market, prices should correspond to the intrinsic value. This deviation decreased from an average of 1.36 in the first ten rounds to 0.84 in the second half of the experiment. This growing efficiency

[13] Due to the zero-sum property of the market the traders I6 to I9 all had a lower return in the second half of the experiment, but their results are not displayed here for the sake of clarity.

of the market can be attributed to the lower mispricings due to fewer average informed traders relying on their information.

Above we saw, that well informed traders used their information to a larger extent than less informed traders. Now we want to investigate whether more active information processing would have benefited the traders. To explore this, we used the real trading data for an extended simulation.
Instead of the actual bid of trader Ix we inserted the expected bid with perfectly active information processing in the trading sheet, leaving the other nine bids as they were. With this data we calculated profits and losses for all periods for each trader individually. We then compared the profits/losses for trader Ix in our simulation with the real profits/losses in the experiment and checked whether the participant would have been better or worse off with the perfectly active information processing strategy. The net changes for each trader in each of the seven sessions are shown in Table 5. I0 did not have any information to process and is therefore not considered here.

Table 5

Change of profits when using a perfectly active information processing strategy instead of the actual bids

Information level	I1	I2	I3	I4	I5	I6	I7	I8	I9
Session 1	0.46	0.32	0.18	0.48	0.64	0.36	0.10	0.36	0.20
Session 2	-0.12	-0.28	-0.28	-0.25	0.43	0.17	0.31	0.19	0.12
Session 3	0.25	0.13	-0.25	0.00	-0.07	0.15	0.19	0.30	0.00
Session 4	-0.15	0.03	0.13	0.14	0.12	0.00	0.29	0.26	0.09
Session 5	0.23	0.17	0.24	-0.17	0.28	0.25	0.19	-0.01	0.04
Session 6	-0.26	-0.17	-0.13	0.09	0.18	0.19	0.59	-0.01	0.13
Session 7	-0.03	-0.04	-0.12	0.17	0.06	0.00	0.46	0.11	0.20
Average	**0.05**	**0.02**	**-0.03**	**0.07**	**0.23**	**0.16**	**0.30**	**0.17**	**0.11**

In session 2, for example, the trader with information level I1 would have reduced his profits, on average, by 12 Eurocent in each trading period if he had processed his information in a perfectly active way (given all other real bids had not changed). It is noteworthy that traders with information levels I1 to I4 would have improved their return just as often as they would have reduced it. In the aggregate, it makes no difference for them whether they use their information or not. This experimental result is consistent with the equilibrium solution in the simulation, where the marginal returns of both strategies were almost the same.

The situation is different for well informed traders with information levels I5 to I9. They would have increased their profits almost always by switching to the active information processing strategy.[14] Again, we find that there is no single "perfect strategy" on how to process information in a market. It rather depends on the trader's existing information level and the actions of all other traders.

6 Concluding Remarks

In this paper we explored a field, yet little covered in finance – the relationship of information level and return in markets – by using ten different information levels in a virtual market with heterogeneously informed agents. In addition to the "usual" result of informed traders outperforming uninformed ones, we find an area where additional information is of no value or even harmful. We understand the financial market as a game, where heterogeneous agents interact. In their attempts to outsmart each other only the best informed are able to gain above average returns. Average informed traders who rely on their information can be exploited by their better informed opponents, while random traders can expect the average return, as they are not exploitable. This relationship was tested in a simulation and an experiment.

By changing trading strategies of underperforming agents in our simulation we derive an equilibrium where no trader can improve his situation by changing his trading strategy. In this equilibrium well informed traders process their information actively, while low to average informed traders act randomly. The marginal return of the different trading strategies is approximately the same.

In this simulation we also find a decreasing marginal return of trading strategies, what we consider one of the most remarkable results of our study. If more traders use a strategy, the results with this strategy will deteriorate. A general trading rule for low and average informed agents one could derive from this is to use a different trading strategy, than the others ("never follow the herd").

The simulation results were confirmed in an experiment using the same market design to observe the behaviour of human beings in this environment. We found that the at

[14] 32 of the 35 traders would have improved their profits by up to 64 Eurocent per period, while just three traders would have been worse off by zero to seven Eurocent. The distribution of the number of winners and losers when comparing the actual bid with the expected bid for perfectly active information processing is significantly different between the traders with information levels I0 to I4 and I5 to I9 respectively ($p < 0.01$, χ^2-test, $df = 1$).

least some participants were able to learn from their mistakes during the experiment and to correct them to a certain extent. This change in trading patterns led to average returns that resembled the equilibrium we got in the simulation approach. If the traders had used their information more actively, the best informed would have profited even more, while it makes basically no difference for low to average informed traders whether they use their information or not.

Further research in this field promises to be fruitful. We plan to continue by introducing an information market to auction or sell the information to participants. The reasons and mechanisms for individuals to gather information in a market is one more mystery yet unsolved.

References

Ackert, L., Church, B., Zhang, P., 2002. Market behaviour in the presence of divergent and imperfect private information: experimental evidence from Canada, China and the United States. Journal of Economic Behavior & Organization, Vol. 47, p. 435-450.

Arthur, B., Holland, J., LeBaron, B., Palmer, R., Tayler, P., 1997. Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. In The Economy as an Evolving Complex System II, Eds.: W. Brian Arthur, Steven N. Durlauf, David A. Lane, Boulder, Colorado, Westview Pr.

Blackwell, D., 1951. Comparison of experiments. In J. Neyman (ed.), Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, Berkeley, University of California Press.

Copeland, T., Friedman, D., 1992. The Market Value of Information: Experimental Results. Journal of Business, Vol. 65, p. 241-265.

Diamond, D., Verrecchia, R., 1981. Information aggregation in a noisy rational expectations economy. Journal of Financial Economics, Vol. 9, p. 221-235.

Fama, E., 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. Journal of Finance, Vol. 45, p. 383-417.

Fischbacher, U., 1999. Z-tree: Zürich Toolbox for Readymade Economic Experiments. Working paper No. 21, Institute for Empirical Research in Economics, University of Zurich.

Gibbons, R., 1992. A Primer in Game Theory, Prentice Hall, Gosport.

Grossman, S.J., 1976. On the Efficiency of Competitive Stock Markets where Traders have Diverse Information. Journal of Finance, Vol. 51, p. 573-585.

Grossman, S.J., Stiglitz, J.E., 1980. On the Impossibility of Informationally Efficient Prices. American Economic Review, Vol. 70, p. 393-408.

Haltiwanger, J., Waldman, M., 1985. Rational Expectations and the Limits of Rationality: An Analysis of Heterogeneity. American Economic Review, Vol. 75 (3), p. 326-340.

Hirshleifer, J., 1971. The Private and Social Value of Information and the Reward to Inventive Activity. American Economic Review, Vol. 61, p. 561-574.

Malkiel, B.G., 2003a. Passive Investment Strategies and Efficient Markets. European Financial Management, Vol. 9 (1), p. 1-10.

Malkiel, B.G., 2003b. The Efficient Market Hypothesis and Its Critics. Journal of Economic Perspectives, Vol. 17 (1), p. 59-82.

Schredelseker, K., 1984. Anlagestrategie und Informationsnutzen am Aktienmarkt. Schmalenbachs Zeitschrift für Betriebswirtschaftliche Forschung, Vol. 36/1, p. 44-59.

Schredelseker, K., 2001. Is the usefulness approach useful? Some reflections on the utility of public information. In Contemporary Issues in Accounting Regulation, McLeay, Stuart and Riccaboni, Angelo (Eds.), p. 135-153.

Strong, N., Walker, M., 1987. Information and Capital Markets. Worcester, Basil Blackwell.

Sunder, S., 1992. Market for Information: Experimental Evidence. Econometrica, Vol. 60, p. 667-695.