On the marginal benefits of additional information in markets with heterogenously informed agents - an experimental study[1]

Jürgen Huber[2], Michael Kirchler and Matthias Sutter
University of Innsbruck

March 9, 2004

[1]We are grateful for the financial support of the Raiffeisen-Landesbank Tirol and of the Center of Experimental Economics of the Institute for Public Economics at the University of Innsbruck.

[2]Corresponding author's address: Institut für Betriebliche Finanzwirtschaft; Universität Innsbruck; Universitätsstraße 15; A-6020 Innsbruck; Tel: 0043 512 507 7554; e-mail: juergen.huber@uibk.ac.at.

Abstract

We examine stock market traders? marginal benefits of additional information on the intrinsic value of a stock market security. We use the method of experimental economics to control carefully for the degree of traders? information. Contrary to conventional wisdom we find that it is possible that a marginal unit of additional information does not have a positive marginal benefit for a trader?s profits. Relatively bad informed traders can even lose money by using their (limited) available information. However, well informed traders benefit significantly from more information and from using their information when trading.

1 Introduction

The marginal benefit of additional information is, in general, considered to be positive. 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 (or, more precisely, having it known to the other players that one has more information) can make a player worse off? (Gibbons, 1992, p. 63). In this paper we report the results of an economic experiment, where we examine the marginal benefits of additional information for the profit/loss in a financial market. In our market, we have ten traders, who trade a (virtual) paper, with each trader having a different level of information about the intrinsic value of the paper. Holding all other conditions constant, we can study the marginal benefits of additional information for profits by varying the information level traders, i.e. their information about the ?true? value of the paper. Our results show that additional information only helps to increase the profits of already well informed traders. Traders with low or medium information can even suffer from negative marginal benefits of getting additional information. The rest of the paper is organized as follows: Section 2 gives a brief overview of related literature. The experimental design is introduced in Section 3. The results of the experiment are presented in Section 4, before concluding in Section 5.

2 On the marginal benefit of information

Until the 1960ies the assumption of a positive marginal benefit of information was unchallenged. This changed dramatically in 1970, when Fama (1970) published his famous paper on the information efficiency of markets. The efficient market hypothesis (EMH) simply states, that ?prices fully reflect all available information.? (Fama, 1970, p. 385). However, if prices react to new information instantly and fully reflect all available information, one should not observe a positive relation between a trader?s level of information and his profits from trading (Hirshleifer, 1971). The expected return from trading would be the average market return, no matter how well the trader is informed. The EMH in its strong form has soon been challenged both by empirical evidence of so-called ?anomalies? (calendar effects, size effects, etc.; see, for instance, Shiller, 2003) as well as by theoretical arguments. Grossman (1976) introduced the ?information paradox? associated with the EMH, meaning that if prices always fully reflect all information, then there is no incentive for actors to gather (costly) information, as this could not improve their return. But if nobody gathers information, there is no reason to assume that prices reflect all ? or even any ? information. In this situation it would clearly pay off to be better informed, which contradicts the EMH.1 Grossman and Stiglitz (1980) solve this problem by assuming a positive marginal return from additional information which is equivalent to the marginal costs of gathering additional information. Therefore, the net return ? taking into account information gathering costs ? would again be the same for all investors, irrespective of their information level. Today the strongest form of the EMH is only considered a theoretical concept, but not a valid model for real markets. In a survey on the EMH and its development, Fama wrote in 1991: ?Since there are surely positive information and trading costs, the extreme version of the market efficiency hypothesis is surely false.? (Fama, 1991, p. 1575). This is underlined by empirical studies showing that insiders are able to gain above-average returns (see Jeng et al., 2003, and the references cited there). The central concern of our paper is to examine whether additional information always has a positive marginal benefit for a trader's profits from trading (i.e. without taking into account information gathering costs). Schredelseker (1984a, 1984b) has argued that negative marginal benefits of additional information are feasible. His line of reasoning runs as follows: A completely uninformed trader can expect the earn the average market return, as follows form the random walk hypothesis. Insiders receive an above-average return, which implies that some traders with an average information level have to earn below-average returns. Given this is true, the relationship between the level of information and the return cannot be monotonically positive, but must have parts with a negative slope as well, indicating negative marginal benefits of additional information. Evidence on 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, 2003). To the best of our knowledge the possibility of a negative marginal benefit of additional information has not been subject to a rigorous test. Schredelseker (2002) has run several simulations with computerized traders who have different levels of information about the ?true? value of a (hypothetical) paper. He has found that low-to-medium informed investors suffer from negative marginal benefits of using additional information, while gathering additional information paid off for already well informed traders. However, the problem with simulations is their results are predetermined by the algorithm used for simulation traders? decisions. Furthermore, traders are not real persons suffering losses or making profits from their investment decisions, but rather computers who face no consequences of their decisions. The investigation of the marginal benefit of additional information with field data (e.g. with traders in real stock exchanges) would be most desirable, but faces several problems: The information level of the traders can not be controlled or judged adequately, because it is unclear what information is relevant to which extent. The ?true? value of a stock is never known, which makes classifying trading decisions as good or bad very complicated. Yet, the positive marginal benefit of insider information is empirically affirmed (Jeng et al., 2003), but field data on low and medium informed traders will hardly ever be sufficiently quantifiable to gain reliable results. Moreover, field data nearly always violate the ceteris paribus condition. When market situations change rapidly ? as they do (think of, e.g., volatile prices, trading volumes, etc.) ? the influence of the information level on total return can not be controlled for adequately, even if the information level of traders could be specified correctly. With laboratory experiments information levels can be varied and controlled systematically while keeping other variables constant (ceteris paribus). This methodological advantage has led to numerous studies examining the influence of information on market prices and the return of individual traders in these markets (Plott and Sunder, 1988; Copeland and Friedman, 1992; Sunder, 1992; Sunder, 1995, provides a survey). Most studies were interested in the value of insider information, which was basically confirmed as being very valuable (see, e.g., Güth et al., 1997). However, all studies consider only two types of traders: Informed and uninformed ones. Given this binary informational structure, it is not possible to examine whether the marginal benefit of additional information is always positive, even when several different information levels exist. Our experimental design is adequate to examine this question.

3 The Experiment

3.1 Design

Our experimental design is based on the simulation approach chosen by Schredelseker (2002), but we have real subjects who can earn real profits from trading. In our market ten subjects trade a (virtual) paper in 20 independent rounds. In each round the intrinsic (or ?true?) value of the paper is given by the sum of ten Laplace-coins showing either ?1? or ?0?. The ten different coins can be seen as ten different kinds of information (e.g. one coin being the general economic outlook, another coin being the prospects on inflation, another one being the outlook on company research projects, etc.), with ?1? indicating ?good? and ?0? indicating ?bad?. Each round each trader has to make a bid for the paper. The median of these bids becomes the market price, with all traders having posted a lower bid becoming sellers (short) and all traders bidding more becoming buyers (long). Traders bidding exactly the market price are neutral in the respective round and have neither profits nor losses. 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 have been neutral. At the start of the experiment each trader gets a different information level which is fixed for the whole experiment (in order to be able to observe possible learning effects). We have one trader knowing nothing, the next knowing the realization of the first coin, the third knowing coins 1 and 2, etc. until the last trader knowing nine of the ten coins. In the following, we will denote by ?Ix? the information level of trader x, with x ? 0, 1, ?, 9. ?I4?, for instance, denotes the trader who knows the states of the first four coins. Nine traders were students at the University of Innsbruck, while I0 was a computer-simulated trader choosing 0 or 10 with a probablility of 0.5 each. This was done, because we did not want to frustrate a trader by knowing nothing for 20 periods. In each round profits and losses are calculated by comparing the market price with the intrinsic value of the paper. A buyer makes a profit if the intrinsic value is higher than the market price, because he bought something valuable relatively cheap. If the intrinsic value is below the market price, a buyer makes a loss. A seller has a profit, 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. To assure that the zero-sum-game-property of the market holds, scale selling is used, if the number of buyers and sellers is not equal.

3.2 Simulation results

To gain a benchmark for our experimental results we ran a set of simulations, using our experimental design, but with computer-generated, rather than human traders. Nine of the ten traders used an active information processing strategy by forming the expected value of the coins they do not know and adding the values of the coins they see, while I0 was a random player as in the experiment. To increase the reliability of the results we ran ten simulations with the complete set of the 1024 possible sets of coins. As can be seen in Figure 1 below, the results confirm the thesis of an area of negative value of additional information.

Figure 1: Profit per period and information level in the simulation

The key to solving this puzzle is biased information: When the distribution of coins is biased in one direction (e.g. the first coins showing a very different picture than the others, as 0000111111), the average informed traders are all mislead in the same direction (here: Estimating the intrinsic value too low), while uninformed traders know too little to be mislead, and the best informed know enough to make a good estimate. In such cases the average informed will systematically loose money, while the others win. A further discussion of this effect will follow when we talk about the results of the experiment.

3.3 The determination of the intrinsic value

In realising the ten sequential coin flips of a trading period we have 1024 (=210) different possible outcomes. To get comparable conditions in all experimental sessions we have decided not let the coins flip randomly, but to pre-select 20 out of the 1024 possible coin sequences. This procedure ensures comparability across experimental sessions and is necessary for ?clean? statistical tests. We have chosen 20 coin sequences (of 10 flips) which are a good representation of the whole set of 1024 possible outcomes. The distribution of the intrinsic values of our 20 sequences is as close as possible to the respective distribution for all possible outcomes, as can be seen from Table 1. The average intrinsic value of our 20 sequences is 5.05. The relative frequency of a given intrinsic value, ranging from 3 to 6, differs less than 2 percentage points between our 20 sequences and the whole set of 1024 possibilities. For the other intrinsic values, the maximum differences between our 20 sequences and the whole set is 4.4 percentage points, which seems reasonable close to the whole set. Our set of 20 sequences (of ten coins) have been randomly ordered and kept constant in all sessions to avoid order effects.

Table 1: Intrinsic value of the experimental coins sets and of all 2^10 coin flips

Intrinsic value	0	1	2	3	4	5	6	7	8	9	10
total frequency											
Experiment	0	1	0	2	4	5	4	3	1	0	0
Whole set	1	10	45	120	210	252	210	120	45	10	1
relative frequency (%)											
Experiment	0,0	5,0	0,0	10,0	20,0	25,0	20,0	15,0	5,0	0,0	0,0
Whole set	0,1	1,0	4,4	11,7	20,5	24,6	20,5	11,7	4,4	1,0	0,1

Experiment: N = 20 periods with 10 coins each Whole set: N = 1024 different flips with 10 coins each

3.4 Implementation of the experiment

The experiment was run computerized (with the help of z-Tree, Fischbacher, 1999) on June 19th 2002 with seven groups of nine students each at the University of Innsbruck. On average, sessions lasted about 60 minutes and the average payment to the students was ? 14. In each round the participants had to enter their bid for the paper on a computer screen. In addition they got a set of information including the states of the coins (?0? or ?1?) for their respective information level (I0 to I9). In the lower part of the screen particants could trace their decisions and the market outcome in all previous rounds. In particular, they could see their bid, the market price, their position on the market (buyer/seller/neutral), the intrinsic value of the paper, the profit or loss for a given round and their total profit.2

4 Experimental results

4.1 Information level and profit

The most obvious research question is the relationship between the information level of the participants and their return. Figure 2 shows the average profit per period for each of the ten information levels. We see that up to I5 all six traders make losses which are roughly the same, with I0 loosing least and I4 loosing most. It is remarkable that we do not find increasing returns for better informed traders in the range from I0 to I5, but even slightly negative returns. However, the differences between the return of I0 to I5 are statistically not significant (Friedman two-way analysis of variance). We can conclude that in this area of low to medium informed traders the marginal benefit of additional information is basically zero. The deviation of theses results from the simulation can be explained by the trading strategies employed by the participants, which will be the topic of the next section.

Figure 2: Profit per period and information level in the experiment

Only for I6 and higher information level we find a positive marginal benefit of additional information. Form I6 to I9 the average return increases from 10 eurocent to 52 eurocent per period. The increase in average profits is statistically significant from I5 to I6, I7 to I8 and I8 to I9 (p ¡ 0.05; two-sided Wilcoxon signed-ranks test; N = 7). This finding suggests that additional information is only useful for already well informed traders.

4.2 Information processing strategies in the experiment

In the following we take a closer look at how traders with different information levels use their available information. In a first step we calculate the correlation coefficient according to Pearson (rxy) between the actual bid and the (theoretical) bid with a perfectly active information processing strategy. The latter is calculated by summing up all the known coins and adding the expected value of the unknown coins ? i.e. 0,5 times the unknown coins.3 The higher this correlation coefficient, the higher is the degree of active information processing by the actors, meaning that the trader uses the information available to him. In the upper part of Table 2 we see the respective correlation coefficients for all 20 periods of the experiment for each of the ten information levels. We find that the correlation increases systematically with the information level. The relatively well informed traders knowing six to nine coins use their information heavily when forming their bid. The least informed traders knowing zero or one coin basically ignore even the little information provided to them. I1 for example has an expected bid of 4.5 (if the first coin shows 0) or 5.5 (if the first coin shows 1) for each period, but we found, that most traders chose a random strategy. For the average informed traders we find a correlation around 0,4. These traders use their information most of the time, but sometimes they ignore it as well.

Table 2: Information processing in the experiment Information level I1 I2 I3 I4 I5 I6 I7 I8 I9 Correlation between bid and expected bid with active information processing

-0,03

0,38

0,40

0,40

0,58

0,81

0,78

0,74

0,85 Mean average deviation of bid from the expected bid with active information processing

1,39

0,81

1,23

1,01

0,88

0,55

0,65

0,85

0,70

In the lower part of Table 2 we used an additional indicator for the degree of information processing. We calculate the mean average deviation of the actual bid from the expected bid with a perfectly active information processing strategy. We find that the mean average deviation decreases with the information level. This can be interpreted as further evidence that the well informed traders use their information more actively than the less informed traders.

4.3 The benefit of perfectly active information processing

In the previous section we saw, that well informed traders used their information to a larger extent than less informed traders. Now we want to look 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 a trader with information level Ix we inserted the expected bid with perfectly active information strategy in the trading sheet, leaving all the other nine bids as they were. With this data we calculated profits and losses for all periods. We then compared the profits/losses for trader Ix in our simulation with the real profits/losses in the experiment and checked whether trader with information level Ix would have been better or worse off with the perfectly active information processing strategy. The net change for each of the seven groups are shown in Table 3. I0 did not have any information to process and is therefore not considered here.

Table 3: Change of profits when using a perfectly active information processing strategy instead of the actual bids used Information level I1 I2 I3 I4 I5 I6 I7 I8 I9 Group 1 0,46 0,32 0,18 0,48 0,64 0,36 0,10 0,36 0,20 Group 2 -0,12 -0,28 -0,28 -0,25 0,43 0,17 0,31 0,19 0,12 Group 3 0,25 0,13 -0,25 0,00 -0,07 0,15 0,19 0,30 0,00 Group 4 -0,15 0,03 0,13 0,14 0,12 0,00 0,29 0,26 0,09 Group 5 0,23 0,17 0,24 -0,17 0,28 0,25 0,19 -0,00 0,04 Group 6 -0,26 -0,17 -0,13 0,09 0,18 0,19 0,59 -0,01 0,13 Group 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 group 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 increased just as often as they would have reduced their profits. In the aggregate, it makes no difference for them whether they use their information or not. However, the situation is completely 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.4 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.

4.4 Learning effects

So far, we have seen that traders with information levels I0 to I5 lose money, on average, throughout the 20 round of the experiment. In Table 4, we show whether the relatively little informed traders were able to reduce their losses in the course of the experiment. To check for the possibility of learning, we consider the profits/losses in the first and the second half of the 20 rounds. I0 is included for completeness, but of course the improvement of this computer player is by chance.

Table 4: Average profits/losses in first vs. second half of the experiment Information level I0 I1 I2 I3 I4 I5 I6 I7 I8 I9 1) Rounds 1 ? 10 -0,15 -0,22 -0,19 -0,22 -0,47 -0,19 0,19 0,16 0,44 0,65 2) Rounds 11 ? 20 -0,13 -0,18 -0,13 -0,13 0,05 -0,21 0,00 0,10 0,24 0,39 Difference 2) ? 1) 0,02 0,04 0,05 0,09 0,52 -0,02 -0,19 -0,06 -0,20 -0,26

It is obvious that the low to medium informed traders (I0 to I4) were able to reduce their losses, with a significant increase of profits for traders with information level I4 (p ¡ 0.01; two-sided Wilcoxon signed-ranks test; N = 7). Conversely, the better informed traders (I5 to I9) had lower profits in the second half of the experiment. The latter result is, of course, driven by the zero-sum-property of the market. The improvement in the average profits of traders with information levels from I0 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.

5 Conclusion

Our experimental results show that additional information need not necessarily have a positive marginal benefit for a trader?s profits in a market. Especially for relatively bad informed traders the marginal benefit of additional information is basically equal to zero, in our experiment it is even slightly negative, on average. Most of these traders would profit by simply ignoring the (little) information they have. This can be explained quite easily: In a market traders basically make bets on future developments of prices. Bad informed trader always make bets against ?smarter? people, which typically leads to losses. Choosing (or bidding) randomly may actually improve the performance of badly informed traders. This is, basically, what we observe for low and medium information levels in the second part of our experiment. In our study we disregarded information costs to make the framework for our analysis as simple as possible. Under the (realistic) assumption of positive information costs, the net profits (after deduction of information costs) would actually fall with an increase in traders? information level from I0 to I5. We are aware that this reasoning is a challenge to the thesis of Grossman and Stiglitz (1980), who assumed that the expected returns in a market would increase with the information level to cover information costs. Our results, however, imply that there are non-linear effects in the marginal benefits of additional information. Only already well informed traders should gather further information, while bad informed traders better ignored even the little information they have. If this conclusion is accepted, then we should find only very well or very badly informed traders in a market, but very few traders with average information. Future research might address this issue by including a market for gathering information in our trading market.

References * Copeland, Thomas and Friedman, Daniel (1992), The Market Value of Information: Experimental Results, in: Journal of Business, Vol. 65, p. 241-265. * Diamond, Douglas W. and Verrecchia, Robert E. (1981), Information Aggregation in a Noisy Rational Expectations Economy, in: Journal of Finance, Vol. 56, p. 221-235. * Fama, Eugene (1970), Efficient Capital Markets: A Review of Theory and Empirical Work, Journal of Finance, Vol. 45, p. 383-417. * Fama, Eugene (1991), Efficient Capital Markets II, Journal of Finance, Vol. 66, p. 1575-1616. * Fischbacher, Urs (1999), Z-tree: Zurich Toolbox for Readymade Economic Experiments. Working paper No. 21, Institute for Empirical Research in Economics, University of Zurich. * Gibbons, Robert (1992), A Primer in Game Theory, Prentice Hall, Gosport. * Grossman, Sanford J. (1976), On the Efficiency of Competitive Stock Markets where Traders have Diverse Information, Journal of Finance, Vol. 51, p. 573-585. * Grossman, Sanford J. and Stiglitz, Joseph E. (1980), On the Impossibility of Informationally Efficient Prices, American Economic Review, Vol. 70, p. 393-408. * Güth, Werner, Krahnen, Jan P. and Rieck, Christian (1997), Financial Markets with Asymmetric Information: A Pilot Study Focusing on Insider Advantages, Journal of Economic Psychology, Vol. 18, p. 235-257. * Hirshleifer, Jack (1971), The Private and Social Value of Information and the Reward to Inventive Activity, American Economic Review, Vol. 61, p. 561-574. * Jeng, Leslie A., Metrick, Andrew and Zeckhauser, Richard (2003), Estimating the Returns to Insider Trading: A Performance-Evaluation Perspective, The Review of Economics and Statistics, forthcoming. * Kyle, Albert S. (1989), Informed Speculation with Imperfect Competition, Review of Economic Studies, Vol. 57, p. 317-356. * Malkiel, Burton G. (2003), The Efficient Market Hypothesis and Its Critics, Journal of Economic Perspectives, Vol. 17 (1), p. 59-82 * Plott, Charles and Sunder Shyam (1988), Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets, Econometrica, Vol. 56, p. 1085-1118. * Schredelseker, Klaus (1984a), Anlagestrategie und Informationsnutzen am Aktienmarkt, Schmalenbachs Zeitschrift für Betriebswirtschaftliche Forschung, Vol. 36, p. 44-59. * Schredelseker, Klaus (1984b), Dummheit schlägt Mittelmäßigkeit - Zur Kritik Bernhard Kromschröders am Informationsnutzenmodell des Aktienmarkts, Schmalenbachs Zeitschrift für Betriebswirtschaftliche

Forschung, Vol. 36, p. 1074-1079. * Schredelseker, Klaus (2002), Grundlagen der Finanzwirtschaft, München, Wien, Oldenbourg Verlag. * Seyhun, Hasan N. (1986), Insiders' Profits, Costs of Trading and Market Efficiency, Journal of Financial Economics, Vol. 16, p. 189-212. * Shiller, Robert J. (2003), From Efficient Markets Theory to Behavioral Finance, Journal of Economic Perspectives, Vol. 17 (1), p. 83-104. * Sunder, Shyam (1992), Market for Information: Experimental Evidence, Econometrica, Vol. 60, p. 667-695. * Sunder, Shyam (1995), Experimental Asset Markets: A Survey, in: Kagel, John and Roth, Alvin E. (Hrsg), Handbook of Experimental Economics. Princeton, NJ, Princeton University Press, p. 445-495.

1 Broadly speaking, there are two ways to resolve the problem of the information paradox. First, it is possible to assume some ?noise? in the market, which leads to prices not fully reflecting all information perfectly. In these ?noisy rational expectations equilibrium? models it pays off to gather additional information (Grossman and Stiglitz, 1980; Diamond and Verrecchia, 1981). Second, the assumption of perfect competition can be relaxed. If insiders with market power trade more/less than their information would suggest, information is held back from the market, which can therefore not reflect all information (Kyle, 1989; Jackson, 1991).

2 The starting capital of the participants varied according to their information level - the least informed (I0) got ? 19, while the best informed (I9) started with only ? 6. 3 E.g., using a perfectly active information processing strategy, a trader with I6 seeing 1-1-0-1-0-1 would estimate the intrinsic value to be 4+0,5*4=6. 4 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, ?ý-test, df = 1).