

Cascades in Experimental Asset Markets

Christoph Brunner

September 6, 2010

Abstract

It has been suggested that information cascades might affect prices in financial markets. To test this conjecture, we run experimental double auctions in which subjects receive private signals about the common value of an asset. These signals are released sequentially and subjects are told which other traders already received their signal. We test whether prices in periods in which the first few signals are misleading indeed fail to converge to the expected value of the asset given all private signals but find no significant difference compared to other periods. We also run a control treatment in which subjects receive the same signals simultaneously at the beginning of the trading period but again fail to find significant differences compared to the treatment in which signals are released sequentially. There is some evidence that subjects are trying to mislead other traders by submitting bids and asks that clearly do not correspond to their signals.

1 Introduction

In the mid-1500s, the Netherlands became a center of cultivation and development of new tulip varieties. A market for tulip bulbs was established. From November 1636 to January 1637, prices of rare varieties surged upward and then rapidly collapsed to approximately 10 percent of their peak values (Garber, 1990). This episode, commonly referred to as “tulip mania, is only one example for trade occurring at seemingly irrational prices. Other instances often involve financial markets: in May 1719, the shares of the French Compagnie des Indes sold at approximately 500 livres per share. About 5 months later, the same shares were traded at 10.000 livres. The

surge was followed by a crash: in September 1721, the price was down at 500 livres per share again (Garber, 1990). More recent examples of rapidly rising and shortly afterwards even more rapidly collapsing prices include the stock market movements in the United States at the end of the 1920s or the 1980s. In recent years, the valuations of many Internet-related companies exhibited similar patterns.

Since the expected value of these assets given all available information at the time is unknown, it remains unclear whether the according prices deviated from their fundamental value or not. If they did, one possible explanation is that these were instances of information cascades. According to Bikhchandani et al. (1992), “An informational cascade occurs when it is optimal for an individual, having observed the actions of those ahead of him, to follow the behavior of the preceding individual without regard to his own information” (p. 994). In an information cascade, individually held information is no longer revealed to others. As a consequence, agents can end up making suboptimal choices given all privately held information at the time.

In this chapter, we will examine whether information cascades really routinely occur in double auctions with endogenous prices. After a brief review of the relevant literature, we will discuss the experimental design and then present the according results.

Even though both Bikhchandani et al. (1992) and Banerjee (1992) mention financial markets as a possible environment in which information cascades might occur, agents in their models sequentially choose among a set of options for which they have common values. There are no prices and no trade occurs. In such environments, information cascades are quite frequently observed in the laboratory (Anderson and Holt, 1997, for example). However, it is unclear whether these results are relevant to financial markets. Avery and Zemsky (1998) show that price adjustments prevent information cascades under common knowledge of rationality when there is only 1 dimension of uncertainty. To test whether the presence of prices really eliminates information cascades, Drehmann et al. (2005) and Cipriani and Guarino (2005) implement markets in which the price of the asset always corresponds to the expected value given the history of previous choices. If all agents were fully rational, subjects should always follow their private information in these markets but they often fail to do so. Nevertheless, neither Drehmann et al. (2005) nor Cipriani and Guarino (2005) find evidence for information cascades. Instead, subjects exhibit contrarian tendencies.

Despite these findings, it is not clear that it is impossible for information cascades to be observed in financial markets. Avery and Zemsky (1998) show that information cascades can occur in a model with an uninformed market maker, noise traders and event uncertainty. Agents in their model are only allowed to trade once in an exogenously determined sequence. Lee (1998) relaxes that latter condition and shows that even under common knowledge of rationality, agents are not necessarily always following their private information. However, these are not information cascades in the sense of Bikhchandani et al. (1992) since agents' actions still depend on their private signal, they just might choose not to follow their signal if it is not strong enough.

In the experimental literature, there is some evidence suggesting that information cascades might occur even when prices are endogenous and in the absence of restrictions on the time and quantity of trade. In an experimental double auction, Barner et al. (2005) find that informed traders tend to trade particularly actively at the beginning of the period except in periods in which information aggregation fails. This result suggests that subjects carefully observe what other traders do and update their valuations accordingly. If the first few actions in a trading period are misleading, prices tend not to reflect the expected value of the asset. Plott and Sunder (1982) also report that informed traders are particularly active in the early stages of trading. Other experimental studies find that subjects tend to rely on their private information too often (Nöth and Weber, 1993, for example) or that they purchase too many signals (Krämer et al., 2006). These results would rather suggest that subjects are unwilling to rely too strongly on information revealed by the market and as a consequence, it might be difficult to observe cascades in a market setting.

Hey and Morone (2004) run experimental asset markets in which subjects can purchase signals that indicate the true value of the asset. They report that in 1 period, the first few signals purchased were misleading and as a consequence, prices failed to converge to the true value of the asset. Since they do not report the expected value of the asset given all signals purchased, it is not entirely clear whether information aggregation really failed in this instance. Even if it did, other factors such as a high variance of the signals purchased might have contributed to the failure of prices to converge to the fully revealing rational expectations equilibrium. Moreover, there might have been other periods in which early signals were misleading but information aggregation nevertheless succeeded.

In order to find out whether the sequence in which information is released to the market affects the quality of information aggregation, we design experiments that allow us to control what information is released at which time more closely. Moreover, we run a control treatment in which the same information is released simultaneously. As a result, we can clearly identify periods in which information cascades would occur if agents sequentially chose to buy or sell a unit of the asset at a fixed price. We can then test whether the quality of information aggregation in these periods differs relative to other periods. The treatment in which signals are released simultaneously allows us to control for other possible explanations such as differences in the variance of the signals released.

2 Experimental Design

In order to give subjects an incentive to pay attention to what other agents do, all subjects have the same value for the asset that is being traded. This value is determined independently for each trading period and is equally likely to be 0 or 100 units of the experimental currency. Even in such an environment, subjects often place too much weight on their private information, which would prevent cascades from occurring. One possible way to give traders an even higher incentive to infer what signals other traders received is to release complementary signals. For example, Plott and Sunder (1988) run experiments in which there are 3 possible states of the world, x , y and z . Suppose the true state is y . In that case, half of the traders are told that the state is either x or y while the other half knows that the true state is either y or z . Such an information structure clearly encourages subjects to carefully observe what others do. However, it is not suitable to study cascades since traders would never ignore their private information.

Another possibility is to give more accurate signals to some traders. In that case, it should be more obvious to traders with inaccurate signals that it is in their interest to infer what signals other traders received. Moreover, they are probably more likely to ignore their private information given that others know more about the true value of the asset. Plott (2000) reports convergence of prices to the rational expectations equilibrium in an experiment with heterogeneous private signals.

In order to keep the experimental design as simple as possible, there are only 2 types of signals in our experiments: strong signals and weak

signals. Subjects with weak signals should have every interest in finding out what information subjects with strong signals have. Therefore, the difference between the accuracy of the strong signal and the accuracy of the weak signal should be large. At the same time, even subjects with strong signals should have an interest in observing what others do and as a result, the strong signal should not be too accurate. Also, if weak signals contain almost no information, it would be trivial to find that subjects' decisions do not depend on their private information. For these reasons, the weak signal reflects the true value of the asset with probability 0.6 while the strong signal corresponds to the true value with probability 0.8. Each one of 8 traders is equally likely to receive a strong or a weak signal.

Trading occurs in a continuous double auction that was implemented in jTrade. Each subject is given an endowment of 5 units of the asset at the beginning of each one of 7 trading periods. Subjects also receive a cash loan of 500 units of the experimental currency that they have to repay at the end of the period. Even if prices are at 100, subjects would thus be able to purchase at least 5 units. In order to preserve the symmetry between the buy and the sell side, subjects only have values for at most 10 units of the asset. If a subject ends up holding more than 10 units of the asset at the end of a period, his value for these assets would nevertheless at most be 1000 units of the experimental currency.

In treatment baseline, all subjects receive their private information at the beginning of each trading period. They can then trade for 2.5 minutes. In treatment sequence, the first subject also receives his signal at the beginning of the period and 30 seconds after the market period opens, the second subject receives his private information. After another 30 seconds, another signal is released to the next subject until all 8 subjects received their signals. The position of each subject in this sequence is randomly determined for each period and is shown to other traders along with each bid or ask that the subject submits. As a consequence, all subjects always know whether the trader who submitted a certain bid or ask already received his signal. After all signals are released, subjects are given another 2.5 minutes to trade. As a result, a market period in treatment sequence lasts 6 minutes. This duration is quite different compared to treatment baseline but the time of trading available after all information is released is exactly the same in both treatments. Since we use the same signal and value draws for treatment baseline and treatment sequence, we can test whether adding extra time during which information is released sequentially affects the quality of information aggregation.

Table 1. Experimental Design.

Treatment	# Sessions	# Subjects per Session	Average Earnings
Baseline	5	8	\$20
Sequence	5	8	\$20

We run 5 sessions for each treatment using undergraduate and graduate students at Caltech as subjects. Each subject was allowed to participate only once. In each session, there is 1 practice period that does not affect earnings. Subjects receive a \$5 show-up fee as well as \$1 for every 100 units of the experimental currency. As a result, expected earnings are \$22.5 per subject with sessions typically lasting for 1 hour for treatment baseline and 1.5 hours for treatment sequence.

3 Conjectures and Definitions

Under common knowledge of rationality, no trade should occur in either treatment sequence or treatment baseline (Milgrom and Stokey, 1982). However, we know from previous market experiments that this is rather unlikely to happen. If trade does occur and markets are efficient, prices should always reflect all information available to any trader (Fama, 1970). As a consequence, no information cascades should occur. On the other hand, some subjects might fail to use their private information to update their estimate of the value of the asset and instead rely on what they believe is revealed by actions of other market participants. In that case, information cascades could be observed. More specifically, we examine both “good” and “bad” cascade periods and compare the quality of information aggregation in these periods to other periods.

Definition 1 *A bad (good) cascade period is a period which satisfies the following conditions:*

- *If subjects had to sequentially guess whether the value of the asset was 0 or 100 and the sequence of choices was observable, at least half of*

them would not pay attention to their private signal under common knowledge of rationality.

- *The majority of subjects would choose the wrong (correct) value.*

For the parameters we use in our experiment, only periods in which the first signal is misleading and both the second and the third signal are either weak or both strong and misleading qualify as bad cascade periods. Similarly, only periods in which the first signal is correct and both the second and the third signal are either weak or both strong and correct are good cascade periods. If information cascades are likely to occur in our markets, we expect that the quality of information aggregation in bad cascade periods is worse than in other periods.

Conjecture 1 *The quality of information aggregation for bad cascade periods is higher in treatment baseline than in treatment sequence.*

Conjecture 2 *The quality of information aggregation for bad cascade periods in treatment sequence is lower than for other periods in treatment sequence.*

To measure the quality of information aggregation, we compute the average of the absolute value of the difference between transactions prices and the expected value of the asset given all private signals. We calculate this average for 4 different sets of transactions:

- All transactions that occurred during the last 2.5 minutes of trading
- All transactions that occurred during the last 1.5 minutes of trading
- The last 5 transactions
- The last 5 units that were traded

The last 2.5 minutes of trading are relevant because all information has been released by that time in treatment sequence. Therefore, prices would be identical in both treatments if markets were efficient. We also take the average over transactions that occurred during the last 1.5 minutes of trading in order to allow for time for traders in treatment baseline to reveal their signals to others. For all measures, the mean is always taken over the number

of units that transact. For example, 2 transactions for 1 unit at price x are equivalent to 1 transaction for 2 units at price x . The only exception is the third measure, for which we take the average over the last 5 transactions giving equal weight to each one of them. Transactions at a price of 0 or at prices of 100 or higher are dropped because these are obvious mistakes.

Just like in bad cascade periods, the aggregation of privately held signals stops after some point in good cascade periods if information cascades occur. However, The price at which information aggregation typically stops is almost always quite close to the expected value of the asset given all private signals. Therefore, we do not expect information aggregation in good cascade periods to be substantially worse than in other periods but we nevertheless test whether or not it is. We will also test whether there is a substantial difference in the quality of information aggregation between treatment sequence and treatment baseline using all periods as observations. Since there are only a few bad cascade periods and since we have no reason to expect substantial differences in other periods, we do not expect these differences to be significant.

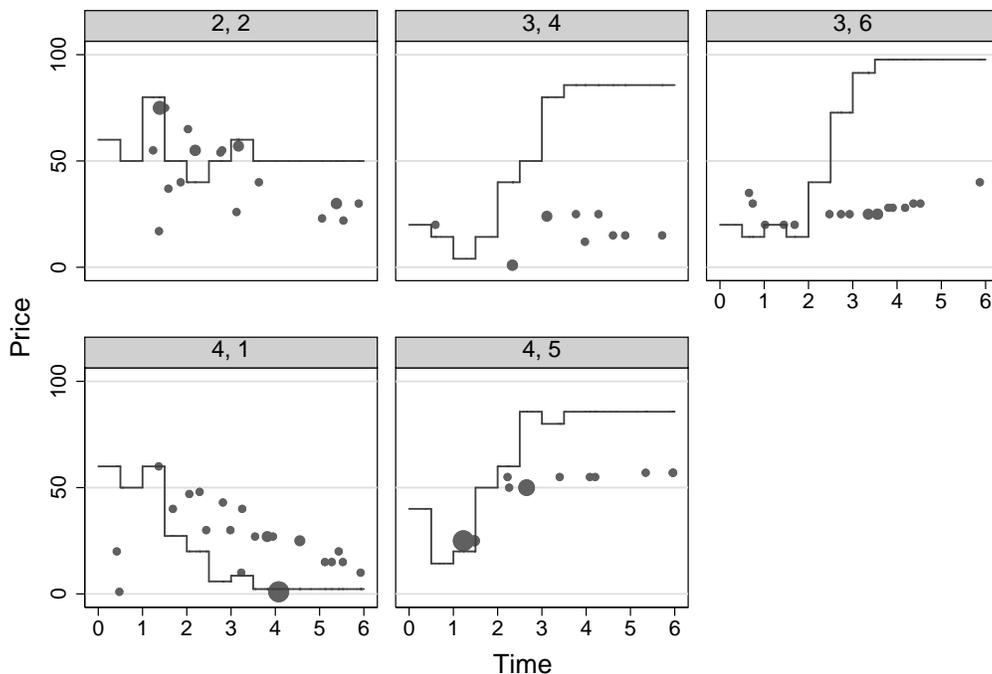
4 Results

In this section, we will first test the conjectures stated above. Since we only find very weak evidence to support them, we then test whether circumstances other than the sequence in which signals are released might be favorable to trigger information cascades. To conclude, we provide some evidence for strategic behavior on the part of subjects, which might explain why we fail to find significant differences between periods that could be favorable to information cascades and other periods in terms of the quality of information aggregation.

4.1 Bad Cascade Periods

To test conjecture 1, we compare bad cascade periods in treatment sequence to bad cascade periods in treatment baseline. Figures 1 and 2 display the according price patterns separately for each one of the 5 periods that qualify as bad cascade periods. The first number at the top of each period-specific graph indicates the session while the second number corresponds to the period. The size of the dots is proportional to the number of units exchanged.

The line corresponds to the expected value of the asset given all private signals. In treatment sequence, prices clearly fail to converge to the rational expectations equilibrium in session 3, both in period 4 and period 6. It could very well be that bad information released early in these periods led to an information cascade that prevented prices from effectively aggregating information. However, prices in the corresponding periods of treatment baseline also failed to converge to the expected value of the asset given all privately held information. As a result, none of the tests we run allows us to reject the null hypothesis that the quality of information aggregation in treatment baseline is equivalent to the quality of information aggregation in treatment sequence for bad cascade periods. Table 2 contains the p-values of Wilcoxon matched-pairs signed-rank tests for the 4 different measures of the quality of information aggregation.



Graphs by wave and period

Figure 1. Prices in Treatment Sequence in Bad Cascade Periods.

Conjecture 2 does not fare much better than conjecture 1. Wilcoxon rank-

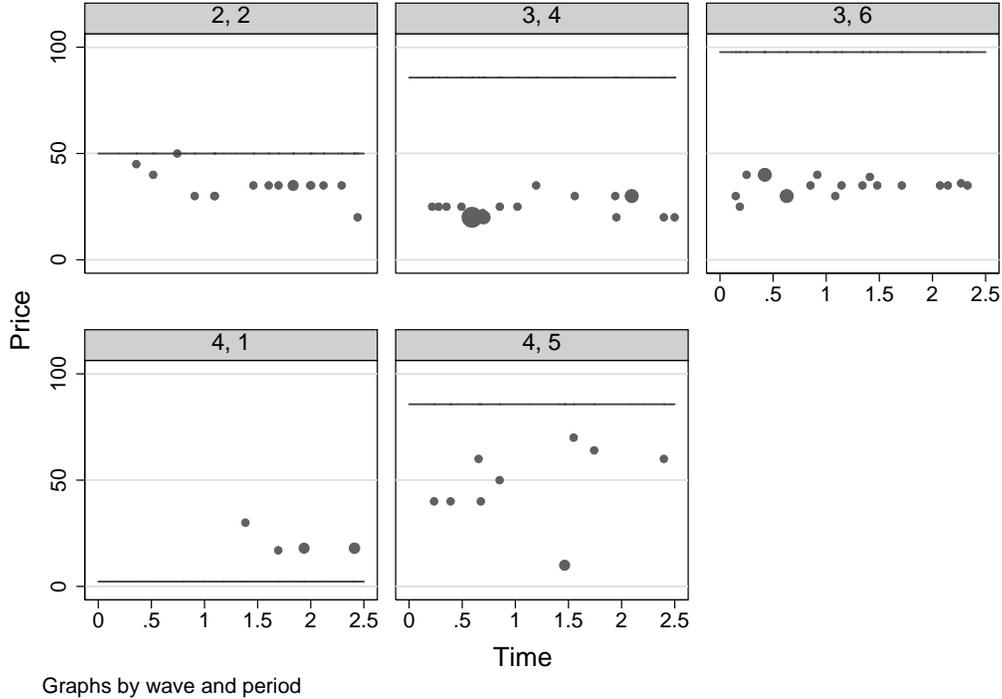


Figure 2. Prices in Treatment Baseline in Bad Cascade Periods.

sum tests do not allow us to reject the null that the median of the quality of information aggregation is identical in bad cascade periods compared to other periods within treatment sequence. When comparing all 35 periods of treatment sequence to all 35 periods of treatment baseline, we are getting fairly close to rejecting the null hypothesis of equal quality of information aggregation for some of the measures employed with treatment sequence exhibiting the higher average absolute deviation of prices from the rational expectations equilibrium price.

The fact that neither one of the main conjectures could be confirmed while there is some support for the hypothesis that the quality of information aggregation is higher in treatment baseline compared to treatment sequence might be due to an insufficient number of bad cascade periods. Clearly, information cascades are not guaranteed to occur even when the sequence of private signals would provide favorable conditions. At the same time,

Table 2. Test Results Bad Cascade Periods.

Measure	Bad Cascade Periods Sequence vs. Bad Cascade Periods Baseline	Bad Cascade Periods Sequence vs. Other Periods Sequence	All Periods Sequence vs. All Periods Baseline
Mean Absolute Deviation Using Transactions During the Last 90 Seconds	Mean Baseline: 39.7 Mean Sequence: 40.1 p = 1.0	Mean Cascade: 40.1 Mean Other: 31.1 p = 0.39	Mean Baseline: 27.6 Mean Sequence: 32.4 p = 0.18
Mean Absolute Deviation Using Transactions During the Last 150 Seconds	Mean Baseline: 39.7 Mean Sequence: 39.9 p = 0.81	Mean Cascade: 39.9 Mean Other: 30.7 p = 0.57	Mean Baseline: 31.6 Mean Sequence: 32.0 p = 0.59
Mean Absolute Deviation Using the Last 5 Transactions	Mean Baseline: 39.0 Mean Sequence: 39.8 p = 1	Mean Cascade: 39.8 Mean Other: 31.5 p = 0.51	Mean Baseline: 27.5 Mean Sequence: 32.7 p = 0.11
Mean Absolute Deviation Using the Last 5 Units Traded	Mean Baseline: 39.8 Mean Sequence: 40.2 p = 0.81	Mean Cascade: 40.2 Mean Other: 31.4 p = 0.45	Mean Baseline: 27.4 Mean Sequence: 32.7 p = 0.14

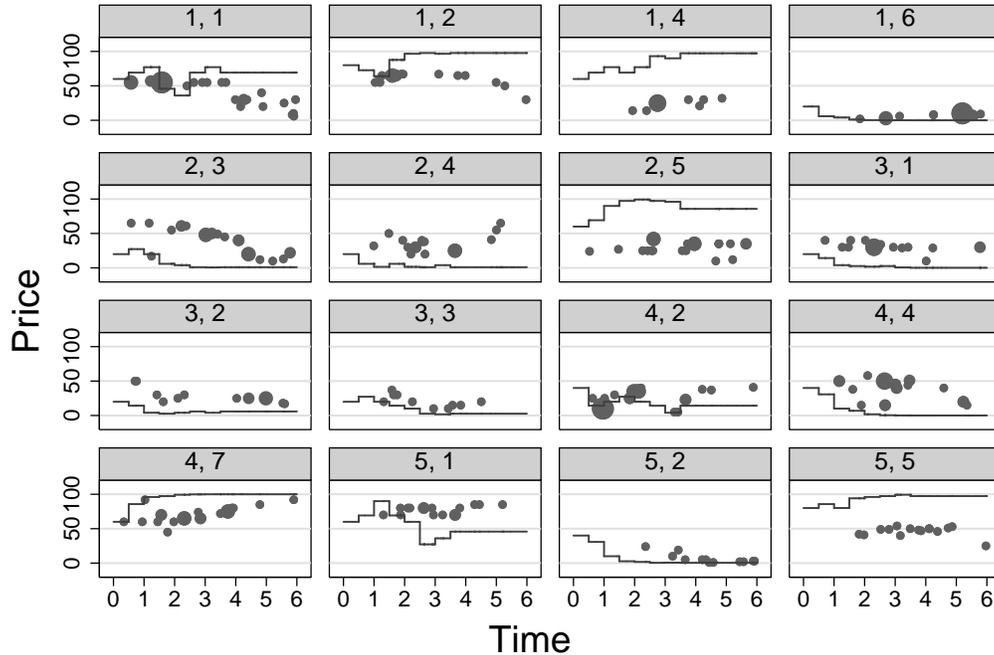
Table 3. Test Results Good Cascade Periods.

Measure	Good Cascade Periods Sequence vs. Good Cascade Periods Baseline	Good Cascade Periods Sequence vs. Other Periods Sequence
Mean Absolute Deviation Using Transactions During the Last 90 Seconds	Mean Baseline: 23.6 Mean Sequence: 32.4 p = 0.11	Mean Cascade: 32.4 Mean Other: 32.4 p = 0.89
Mean Absolute Deviation Using Transactions During the Last 150 Seconds	Mean Baseline: 28.5 Mean Sequence: 30.3 p = 0.68	Mean Cascade: 30.3 Mean Other: 33.4 p = 0.95
Mean Absolute Deviation Using the Last 5 Transactions	Mean Baseline: 23.7 Mean Sequence: 31.3 p = 0.15	Mean Cascade: 31.3 Mean Other: 33.9 p = 0.79
Mean Absolute Deviation Using the Last 5 Units Traded	Mean Baseline: 23.5 Mean Sequence: 30.9 p = 0.16	Mean Cascade: 30.9 Mean Other: 34.2 p = 0.74

information is not always aggregated very efficiently in treatment baseline, either.

4.2 Good Cascade Periods

Figures 3 and 4 display the price patterns in good cascade periods for treatment sequence and treatment baseline. The size of the dots is proportional to the number of units exchanged. The line corresponds to the expected value of the asset given all private signals. Table 3 contains the according test results. We run Wilcoxon matched-pairs signed-rank tests to obtain the p-values for the first column and Wilcoxon rank-sum tests for the second column. While only 5 out of 35 periods qualify as bad cascade periods, 16 qualify as good cascade periods. This might be part of the reason why some of the comparisons between treatment sequence and treatment baseline almost yield significant results. It appears that the quality of information aggregation is somewhat lower in good cascade periods in treatment sequence compared to good cascade periods in treatment baseline.



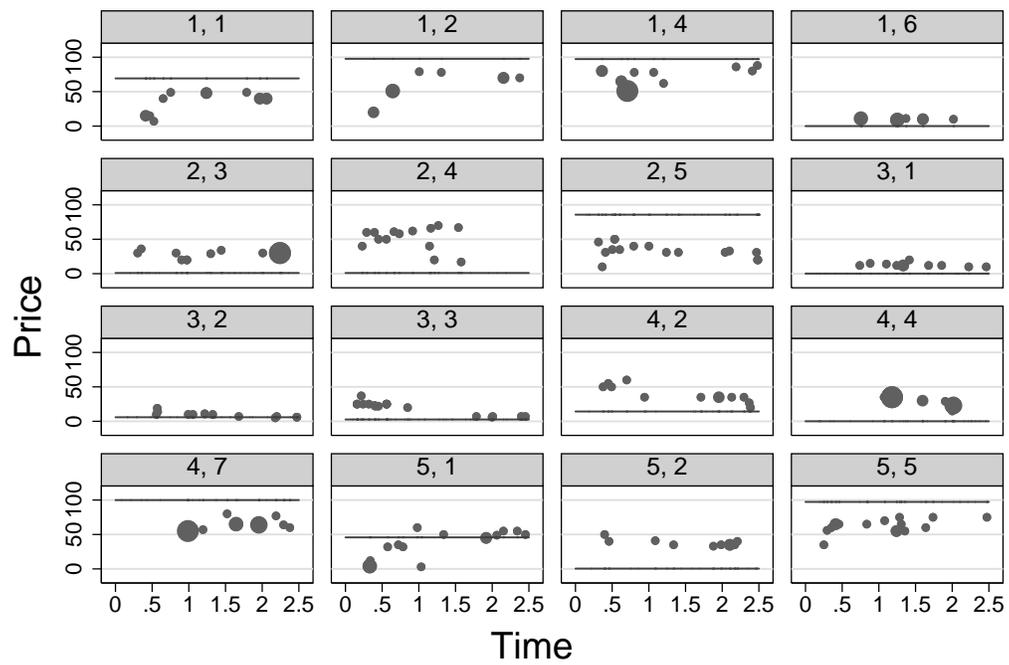
Graphs by wave and period

Figure 3. Prices in Treatment Sequence in Good Cascade Periods.

5 Alternative Cascade Period Definitions

5.1 Shorter Information Cascades

Clearly, the low number of bad cascade periods combined with a relatively high variance in the quality of information aggregation in both treatments contributes to the fact that we did not find much support in favor of conjectures 1 and 2. A possible remedy would be to apply a more liberal definition of cascade periods by relaxing the condition that at least 4 out of 8 traders would ignore their private signal if they chose sequentially whether to buy or sell the asset. However, if only very few traders ignore their private information, it would be difficult to find significant differences with respect to the quality of information aggregation even if an information cascade actually occurred. Therefore, the only extension that we test is one in which at least 3



Graphs by wave and period

Figure 4. Prices in Treatment Baseline in Good Cascade Periods.

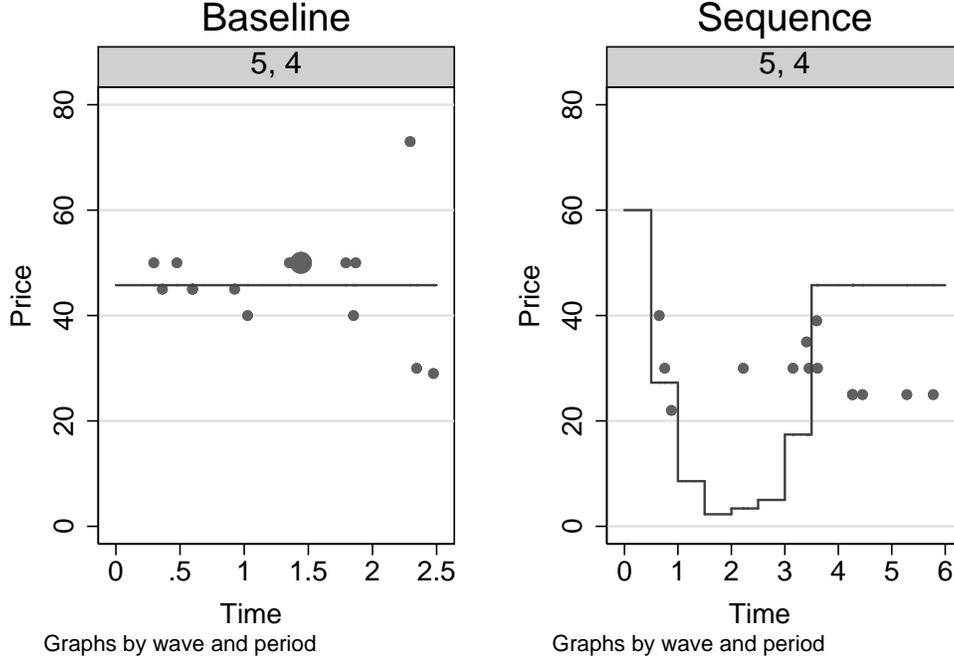


Figure 5. Prices in Session 5, Period 4 for Treatments Baseline and Sequence.

traders would ignore their private information if they chose sequentially. Unfortunately, this more-liberal definition only yields 1 additional bad cascade period (session 5, period 4) and all differences in the quality of information aggregation remain insignificant. Figure 5 displays transaction prices in this additional bad cascade period. The size of the dots is proportional to the number of units exchanged. The line corresponds to the expected value of the asset given all private signals.

5.2 Who Trades First?

Another reason why conjectures 1 and 2 could not be confirmed might be that the sequence in which signals are released does not correspond to the sequence in which subjects actually trade. As a consequence, the sequence in

Table 4. Test Results For Cascade Periods Defined Based on the Time of the First Transaction. We run Wilcoxon rank-sum tests to obtain the p-values.

Measure	Bad Cascade Periods vs. Other Periods Both Treatments	Good Cascade Periods vs. Other Periods Both Treatments
Mean Absolute Deviation Using Transactions During the Last 90 Seconds	Mean Cascade: 31.9 Mean Other: 29.8 p = 0.79	Mean Cascade.: 32.7 Mean Other: 27.4 p = 0.35
Mean Absolute Deviation Using Transactions During the Last 150 Seconds	Mean Cascade: 33.5 Mean Other: 31.6 p = 0.88	Mean Cascade: 35.1 Mean Other: 28.6 p = 0.21
Mean Absolute Deviation Using the Last 5 Transactions	Mean Cascade: 33.8 Mean Other: 29.7 p = 0.57	Mean Cascade: 32.6 Mean Other: 27.7 p = 0.44
Mean Absolute Deviation Using the Last 5 Units Traded	Mean Cascade: 34.9 Mean Other: 29.5 p = 0.45	Mean Cascade: 32.0 Mean Sequence: 28.2 p = 0.55

which private information is revealed to the market might not correspond to the sequence in which signals are released to traders. To test this hypothesis, we compute the Spearman rank correlation coefficient between the position of a subject in the sequence in which private signals are released in treatment sequence and the time at which that subject first buys or sells at least 1 unit of the asset. We obtain a positive correlation (0.24) and can easily reject the null hypothesis of no correlation ($p=0.000$).

Since the correlation is not perfect, we reclassify cascade periods based on the sequence in which subjects trade. We apply the same definition of cascade periods that we originally used (Definition 1) but we now use the sequence of signals obtained by the first 3 subjects who are buying or selling at least 1 unit of the asset to classify periods. In treatment sequence, it occasionally happens that some of these first 3 traders have not yet received their private information at the time at which they first trade. We are not taking these traders into account since their private signal clearly cannot have been revealed to the market at such an early stage. As a result, 7 out of 70 periods qualify as bad cascade periods and 5 of these periods

are from treatment sequence and 3 of them were already originally classified as cascade periods. Wilcoxon rank-sum tests do not allow us to reject the null hypothesis that the median of the quality of information aggregation in these bad cascade periods differs from the median in other periods. When redefining good cascade periods based on the time of the first transaction, 34 out of 70 periods qualify but none of the differences are significant (table 4).

5.3 Who Submits Orders First?

Instead of the time of the first transaction, the time at which a trader first submits a bid or an ask might be more closely related to the time at which he reveals his private information to the market. Therefore, we test whether the time of the first bid or ask is related to the time at which the private signal is received in treatment sequence. The Spearman rank correlation coefficient is 0.29 and we can safely reject the null that the 2 variables are unrelated ($p=0.000$). Instead of taking all bids and asks into account, we only consider bids above 20 and asks below 80 since any signal would justify lower bids or higher asks.

Table 5. Test Results for Cascade Periods Defined Based on the Time of the First Order.

Measure	Bad Cascade Periods vs. Other Periods Both Treatments	Good Cascade Periods vs. Other Periods Both Treatments
Mean Absolute Deviation Using Transactions During the Last 90 Seconds	Mean Cascade: 25.5 Mean Other: 30.0 $p = 0.81$	Mean Cascade.: 31.4 Mean Other: 28.1 $p = 0.82$
Mean Absolute Deviation Using Transactions During the Last 150 Seconds	Mean Cascade: 25.9 Mean Other: 31.1 $p = 0.74$	Mean Cascade: 30.5 Mean Other: 30.9 $p = 0.79$
Mean Absolute Deviation Using the Last 5 Transactions	Mean Cascade: 30.8 Mean Other: 30.0 $p = 0.89$	Mean Cascade: 30.8 Mean Other: 29.5 $p = 0.95$
Mean Absolute Deviation Using the Last 5 Units Traded	Mean Cascade: 31.4 Mean Other: 29.9 $p = 0.79$	Mean Cascade: 30.2 Mean Sequence: 29.9 $p = 0.87$

A reclassification of cascade periods based on the time of the first bid or ask yields 5 bad cascade periods, 4 of these periods already originally qualified as bad cascade periods. We also obtain 35 good cascade periods. Wilcoxon rank-sum tests do not allow us to reject the null hypothesis that the quality of information aggregation does not differ between bad cascade periods and other periods. Good cascade periods also do not seem to yield a different quality of information aggregation (table 5).

6 Alternative Explanations

6.1 Early Expected Value vs. Late Expected Value

Even though subjects might not completely ignore their private information, they might place too much weight on information that the actions of other traders reveal. In that case, misleading early signals would still lead to a lower quality of information aggregation but not necessarily in such a clear-cut way as the cascade model suggests. To measure the extent to which early signals are misleading, we compute the expected value of the asset given the first 4 signals. We then take the absolute value of the difference between this early expected value and the expected value given all private signals. This variable (*devalue*) is then used to explain the quality of information aggregation. Using OLS, we estimate a coefficient for variable *devalue* (β) as well as an intercept (α) separately for treatment sequence and treatment baseline using all 35 periods for each treatment as observations. The results of these regressions are displayed in table 6.

In treatment sequence, the coefficient of *devalue* is always significant at the 10% level. The larger the difference between the early expected value of the asset and the late expected value of the asset, the larger the mean absolute deviation of prices from the late expected value. In that sense, misleading early signals do have an effect on the extent to which prices converge to the rational expectations equilibrium price.

Since *devalue* is correlated with the variance of the signals that traders receive, it could be that a higher variance of signals is the true cause for the observed worse quality of information aggregation in periods with high values of *devalue*. In that case, we would expect the coefficient of *devalue* to be significant in treatment baseline as well. Since that is not the case, we conclude that the sequence in which signals are released does indeed affect

Table 6. Using Devalue to Explain Differences in the Quality of Information Aggregation.

Measure	Treatment	R ²	α (sd)	β (sd)	p-value Wald
Mean Absolute Deviation Using Transactions During the Last 90 Seconds	Baseline	0.06	23.21** (7.735)	0.25 (0.326)	0.30
	Sequence	0.14	24.66** (6.804)	0.44* (0.186)	0.48
Mean Absolute Deviation Using Transactions During the Last 150 Seconds	Baseline	0.07	27.10** (6.754)	0.26 (0.302)	0.27
	Sequence	0.17	23.70** (6.519)	0.47* (0.193)	0.62
Mean Absolute Deviation Using the Last 5 Transactions	Baseline	0.06	22.93** (7.985)	0.26 (0.335)	0.27
	Sequence	0.15	24.61** (6.337)	0.46* (0.191)	0.65
Mean Absolute Deviation Using the Last 5 Units Traded	Baseline	0.07	22.61** (8.103)	0.27 (0.324)	0.29
	Sequence	0.16	24.30** (6.248)	0.48* (0.186)	0.61

Coefficients marked by (* / **) are significant at the (10 / 5) percent level. Robust standard errors clustered by session are shown in parentheses.

the extent to which markets can aggregate privately held information. We also test whether including a variable that measures the standard deviation of signals significantly improves the fit of these regressions. Wald tests do not allow us to reject the null hypothesis that the coefficient of the standard deviation of signals is zero.

6.2 Strategic Behavior

When we classify periods as cascade periods, we assume that subjects reveal their private information to other traders. If they fail to do so, the information that the market receives might not correspond to the information used to classify periods, which could explain why information aggregation in bad cascade periods is not substantially worse than in other periods. In order to test whether subjects are trying to mislead other traders, we examine the first

order submitted in each period. At that time, the only information subjects have is their private signal. We only consider bids above 20 and asks below 80 that were made by traders who had already received their signal. 28% of such first orders are misleading in the sense that subjects are submitting a buy order even though they received a low signal or that they submit a sell order even though they received a high signal. Not all of these orders are inconsistent with the private signal received. For example, if a subject received a weak high signal, the expected value of the asset is 60. It is then perfectly reasonable to submit a sell order at a price of 70. However, 18% of all first orders are either sell orders at a price below the expected value given the trader's private signal or buy orders at a price above the expected value given the trader's signal. Clearly, other traders will find it hard to figure out what signals these traders had based on the bids or asks that they submitted. No matter whether these are intentional attempts to mislead other traders or simply mistakes, the fact that bids and asks do not always reflect the private information that a trader holds makes it difficult to identify the effect of the sequence of signals on the quality of information aggregation.

7 Conclusion

While we find evidence that the sequence in which signals are released to traders affects the quality of information aggregation, there is not much support for conjectures 1 and 2. Bad cascade periods in treatment sequence do not seem to fare substantially worse than bad cascade periods in treatment baseline or other periods in treatment sequence. A possible reason could be that we simply do not have enough observations. Another reason could be the difficulty involved in identifying the sequence in which information is released to the market. In fact, the sequence in which signals are released to traders does not always correspond to the sequence in which subjects actively trade or submit orders. Moreover, when they submit orders, they do not always reveal their private signal but might instead try to mislead other traders.

An alternative experimental design would eliminate these 2 sources of complexity while still preserving an endogenous price. Instead of allowing subjects to trade at any point of time, we could require them to trade in a predetermined sequence. Each subject would be able to submit as many sell and buy orders as desired but only once. As a consequence, subjects

would no longer have an interest in misleading other traders since it would be impossible to capitalize on flawed prices by submitting further orders at a later point of time. At the same time, the sequence in which subjects trade would always correspond to the sequence in which they receive their private information. A control treatment would simply correspond to a call market with identical signals and value draws. By eliminating much of the complexity of a continuous double auction while preserving an endogenous price without a market maker, such a design should allow us to establish whether information cascades routinely occur in markets with an endogenous price.

References

- Anderson, L. R. and C. A. Holt (1997). Information cascades in the laboratory. *American Economic Review* 87, 847–862.
- Avery, C. and P. Zemsky (1998). Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review* 88, 724–748.
- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics* 107(3), 797–817.
- Barner, M., F. Feri, and C. R. Plott (2005). On the microstructure of price determination and information aggregation with sequential and asymmetric information arrival in an experimental asset market. *Annals of Finance* 1, 73–107.
- Bikhchandani, S., D. Hirschleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100(5), 992–1026.
- Cipriani, M. and A. Guarino (2005). Herd behavior in a laboratory financial market. *American Economic Review* 95(5), 1427–1443.
- Drehmann, M., J. Öchssler, and A. Roeder (2005). Herding and contrarian behavior in financial markets — an internet experiment. *American Economic Review* 95(5), 1403–1426.
- Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25, 383–417.

- Garber, P. M. (1990). Famous first bubbles. *Journal of Economic Perspectives* 4(2), 35–54.
- Hey, J. D. and A. Morone (2004). Do markets drive out lemmings—or vice versa? *Economica* 71, 637–659.
- Krämer, C., M. Nöth, and M. Weber (2006). Information aggregation with costly information and random ordering: Experimental evidence. *Journal of Economic Behavior and Organization* 59, 423–432.
- Lee, I. H. (1998). Market crashes and information avalanches. *Review of Economic Studies* 65, 741–759.
- Milgrom, P. and N. Stokey (1982). Information, trade and common knowledge. *Journal of Economic Theory* 26, 17–27.
- Nöth, M. and M. Weber (1993). Information aggregation with random ordering: Cascades and overconfidence. *Economic Journal* 113, 166–189.
- Plott, C. R. (2000). Markets as information gathering tools. *Southern Economic Journal* 67(1), 1–15.
- Plott, C. R. and S. Sunder (1982). Efficiency of experimental security markets with insider information: An application of rational-expectations models. *Journal of Political Economy* 90(4), 663–698.
- Plott, C. R. and S. Sunder (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 56(5), 1085–1118.