

Federal Reserve Bank of New York
Staff Reports

Informational Contagion in the Laboratory

Marco Cipriani
Antonio Guarino
Giovanni Guazzarotti
Federico Tagliati
Sven Fischer

Staff Report No. 715
March 2015



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Informational Contagion in the Laboratory

Marco Cipriani, Antonio Guarino, Giovanni Guazzarotti, Federico Tagliati, and Sven Fischer
Federal Reserve Bank of New York Staff Reports, no. 715

March 2015

JEL classification: C92, G01, G14, G15

Abstract

We study the informational channel of financial contagion in the laboratory. In our experiment, two markets with correlated fundamentals open sequentially. In both markets, subjects receive private information. Subjects in the market opening second also observe the history of trades and prices in the first market. We find that although in both markets private information is only imperfectly aggregated, subjects are able to make correct inferences based on the public information coming from the market that opens first. As a result, we observe financial contagion in the laboratory: Indeed, the correlation between asset prices is very close to that predicted by the theory. Finally, as theory predicts, there is no contagion when asset fundamentals are independent: That is, subjects only react to the history of prices and trades in the first market when it is rational to do so because they convey information.

Key words: information contagion, laboratory experiment

Cipriani: Federal Reserve Bank of New York (e-mail: marco.cipriani@ny.frb.org). Guarino: University College London (e-mail: a.guarino@ucl.ac.uk). Guazzarotti: Bank of Italy (e-mail: giovanni.guazzarotti@bancaditalia.it). Tagliati: University College London (e-mail: federico.tagliati.10@ucl.ac.uk). Fischer: Max Planck Institute for Research on Collective Goods (e-mail: fischer@coll.mpg.de). Guarino gratefully acknowledges the financial support of the Economic and Social Research Council (World Economy and Finance program), the Economic Research Council, and the Institute for New Economic Thinking. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

Comovements among asset prices, whether across countries or across asset classes, are often higher than what can be explained by comovements in asset fundamentals. This empirical regularity, usually referred to as financial contagion, has been widely documented in the empirical finance and international finance literatures.¹ Indeed, almost all the recent episodes of financial turmoil, from the Asian financial crisis of 1997 to the events of 2007-2008, suggest that financial asset prices are very highly correlated, in excess of what could be expected by looking at fundamentals, and that financial instability can quickly spread from one country to the other, or from one asset class to the other.

The theoretical asset pricing literature has highlighted several theoretical mechanisms that can generate contagion in financial markets. In a seminal paper, King and Wadhvani (1990) argued that financial contagion may be explained by informational spillovers across markets. Price changes in a market may be the result of traders reacting to news about an idiosyncratic shock to their market or about a systemic shock. Agents in other markets do not observe whether the shock is idiosyncratic or systemic. They attach some probability to the event that the price movement in the other market is due to a systemic shock, and therefore react to it, even when, in fact, the shock is idiosyncratic. Because of this informational spillover, the correlation among asset prices is higher than that among fundamentals.

The theoretical literature on financial contagion has highlighted other channels generating excess comovements among asset prices in addition to the “information channel” proposed by King and Wadhvani (1990). In Calvo (1999) contagion arises from correlation in liquidity shocks: agents, hit by a liquidity shock in one market, liquidate their position across markets in order to meet a margin call (see also Yuan, 2005). In Kyle and Xiong (2001), financial contagion is due to wealth effects. In Fostel and Geanakoplos (2008) financial contagion arises as a result of the interplay between market incompleteness, agents’ heterogeneity, and margin requirements. In Kodres and Pritsker (2002), contagion happens through cross-market rebalancing, when traders hit by a shock in one market need to rebalance their portfolios of assets. In Pavlova and Rigobon (2007), contagion of stock market prices across

¹See, among the many papers on the topic, Eichengreen *et al.* (1996), Edwards and Rigobon (2002), and Ehrmann *et al.* (2011).

countries arises from wealth transfers and portfolio constraints. The role of information spillovers in generating financial contagion is also illustrated by Cipriani and Guarino (2008): they show that because of informational spillovers across markets, herd behavior and informational cascades transmit from one market to another.²

Informational spillovers across markets are an appealing explanation for financial markets' contagion. For this reason, it is important to test whether the theory accurately describes how learning occurs across markets and causes excess comovements across asset prices. On the one hand, assuming that agents process cross market information rationally (that is, in a Bayesian manner) may underestimate the relevance of informational spillovers. Traders in a market may overreact to what they observe in other markets; for instance, "panic selling" or "irrational exuberance" may occur in a market upon observing a crash or a strong rally in another. On the other hand, traders may focus on their own market and not be able to incorporate all the information content of trading activity in other markets. In other words, it is important to understand whether the theoretical models of information spillover accurately characterize, both in the direction and in the magnitude of the contagion effect, the way traders process information across financial markets.

It is difficult to test the theory's predictions on financial contagion with financial markets data. Financial datasets lack information on several traders' characteristics—in particular, traders' preferences and their information sets—which are unobservable; it is, therefore, hard to test the theoretical model's prediction directly.

To overcome this issue, we brought King and Wadhvani (1990)'s model to the laboratory: in a laboratory financial market, the experimenter directly controls the information set available to subjects acting as traders, and can study the effect of a piece of news regarding one market on other markets. In particular, in the laboratory financial economy that we implemented, traders trade an asset in a standard double auction. Two markets open in sequence, and traders in the second market observe the history of trades and prices in the first. In both markets, traders receive private information about their own asset's fundamental value. Since asset fundamentals

²While the studies just mentioned explain contagion across markets, others have focused on contagion across financial institutions (see, e.g., the seminal contribution of Allen and Gale, 2000; and, for an experimental analysis, Trevino, 2013).

are correlated, information coming from the first market is relevant for the second; and informational spillovers lead to informational contagion across markets.

The results of our experiment show that financial contagion across markets occurs in the laboratory. In fact, the impact of the first market on the second is very close to that predicted by a Rational Expectations Equilibrium. Although private information is not perfectly aggregated in either market, subjects are able to use the public information coming from the market that opens first correctly. As a result, the correlation between asset prices is very close to what the theory predicts. We also studied whether subjects are affected by the history of trades and prices in the first market when theory suggests that they should not. We ran a treatment in which the two asset values were independently distributed. In agreement with the theoretical prediction, we observed no contagion in the laboratory. We conclude that in our study financial contagion does not stem from some irrational reaction to public news, but is the outcome of correct inference by subjects from the public information that they receive from the market opening first.

The rest of the paper is organized as follows. Section 2 presents the theoretical framework. Section 3 illustrates the experiment. Section 4 explains the results. Section 5 concludes. The Appendix contains the instructions of the experiment and additional results.

2 Theoretical framework

2.1 Preliminaries: King and Wadhvani (1990)'s model.

The purpose of our paper is to test the informational contagion channel first proposed by King and Wadhvani (1990) in a controlled laboratory environment. It is therefore useful for us to explain how informational contagion arises in their model.

Let us consider a two-asset financial economy. In both markets, denoted by A and B , the asset price change during a given time interval is a function of newly released information. King and Wadhvani (1990) consider two types of information: i) systemic news—affecting the fundamental values of both assets; and ii) idiosyncratic news, specific to each asset.

For simplicity's sake, consider the case in which the two markets open in sequence, with market A opening first. Traders in market A receive news

about their own asset, and the price of asset A changes accordingly. When market B opens, traders in market B observe the price change occurred in market A , but do not know whether it is due to idiosyncratic or systemic news. As a result, the equilibrium price in market B will change even when the price change in market A was due to purely idiosyncratic reasons. The reason is that traders attach some probability to the event that the price change in market A was due to news about the systemic component, common to both assets.

It is easy to show that the correlation between asset prices in the two markets is higher than what we would observe if traders knew whether a shock is idiosyncratic or systemic (fundamental correlation). King and Wadhwani label this phenomenon “informational contagion”.³

2.2 The model

To test informational contagion in the laboratory, we develop a model that captures the main insights of King and Wadhwani (1990), but is suitable to experimental implementation.

Specifically, we consider a two-market economy, in which the two markets, labeled by A and B , open sequentially. In each market, a continuum of risk-neutral traders trade one asset. The fundamental value of asset A (V^A , which can be thought of as the present discounted value of the asset’s future stream of dividends) takes two values, 0 or 100, with the same probability:

$$V^A = \begin{cases} 0 & \text{with probability } \frac{1}{2}, \\ 100 & \text{with probability } \frac{1}{2}. \end{cases}$$

Although the realization of V^A is unknown to market participants, they have private information about it; in particular, they receive a symmetric binary signal with precision 0.75.⁴ In other words, each participant in market A receives a signal s^A distributed as follows: $\Pr(s^A = 0|V^A = 0) = \Pr(s^A = 100|V^A = 100) = 0.75$.

Market B opens after trading in market A ends. Traders in market B

³King and Wadhwani (1990) contrast this contagion equilibrium, which they label “partially-revealing Rational Expectations Equilibrium,” with a “fully-revealing Rational Expectations Equilibrium” (in which traders observe whether a shock is idiosyncratic or systemic, and although information flows across markets, there is no contagion), and with a “no-communication Rational Expectations Equilibrium” (in which traders do not observe the price in the other market, and no informational spillover occurs across markets).

⁴Obviously, any precision greater than 0.5 would deliver the same qualitative results.

only trade asset B . The fundamental value of asset B , V^B , equals V^A with probability p ; and C with probability $1-p$. C is a random variable distributed as follows:

$$C = \begin{cases} 0 & \text{with probability } \frac{1}{2}, \\ 100 & \text{with probability } \frac{1}{2}. \end{cases}$$

In other words, when $V^A = 0$, then $V^B = 0$ with probability $p + (1-p)\frac{1}{2}$ and $V^B = 100$ with probability $(1-p)\frac{1}{2}$; when $V^A = 100$, then $V^B = 100$ with probability $p + (1-p)\frac{1}{2}$ and $V^B = 0$ with probability $(1-p)\frac{1}{2}$.⁵

Traders in market B receive a symmetric binary signal s^C on the realization of C with precision 0.75, that is, $\Pr(s^C = 0|C = 0) = \Pr(s^C = 100|C = 100) = 0.75$. Furthermore, traders in market B not only observe their own private information, but also the price in market A .

The perfectly competitive Rational Expectations Equilibrium (REE) price of asset A is 0 when $V^A = 0$ and 100 when $V^A = 100$; that of asset B depends both on V^A and on C . If both V^A and C are 0 (or 100), then the price of asset B is 0 (or 100, respectively); if instead V^A and C are different, then the equilibrium price of B is $V^A p + C(1-p)$.⁶

In King and Wadhvani (1990), contagion occurs because agents do not know whether a change in the price of one market stems from an idiosyncratic or systemic shock. In our set up, with probability p the realization of V^A affects V^B , whereas with probability $(1-p)$ it does not;⁷ in the first case, the shock to V^A is systemic (as it affects the fundamental values in both markets), whereas in the second case it is idiosyncratic (and V^B is determined by its own idiosyncratic component). It is easy to show that in the REE the correlation

⁵In King and Wadhvani (1990), the distribution of asset values is

$$\begin{aligned} V^A &= u_A + \alpha u_B + v_A, \\ V^B &= \beta u_A + u_B + v_B, \end{aligned}$$

where u_A , u_B , v_A and v_B are normal random variables and α and β are parameters; u_A and u_B reflect systemic news and v_A and v_B idiosyncratic ones.

In our model, the distribution of asset values is different. Nevertheless, we can interpret the realization of V^A as an idiosyncratic shock to market A with probability $(1-p)$, and as a common shock affecting both markets with probability p ; and the realization of C as an idiosyncratic shock to market B that may occur when the market is not hit by a common shock. Since our assumptions imply that both asset values have the same support $\{0, 100\}$, their distribution is simple to explain to subjects, which makes the model implementation in the laboratory easier.

⁶See Appendix A for the derivation.

⁷To compare our set up with King and Wadhani (1990) we can think of the unconditional value in both markets as being equal to 50 and interpret the realizations of V^A and C as the shock (or the arrival of news) in the market.

between prices is $\frac{p}{\sqrt{(1-2p+2p^2)}}$, whereas the correlation between fundamentals is only p ;⁸ since $\frac{p}{\sqrt{(1-2p+2p^2)}}$ is greater than p , there is contagion from market A to market B .

3 The Experiment

We ran the experiment in the ELSE Experimental Laboratory at the Department of Economics at UCL. Overall, we recruited 192 subjects (undergraduate students in all disciplines) to conduct twelve sessions, six for each of two experimental treatments. Subjects had no previous experience with this experiment and participated in one session only.

3.1 Experimental Design

In order to implement the economy described in Section 2, in each session 16 subjects traded in a two-market, computerized, continuous time, double auction.

We ran two treatments. In the baseline treatment (Treatment I), the value of asset B was equal to the value of asset A with probability 0.5 (i.e., $p = 0.5$); in this case $\Pr(V^B = i|V^A = i) = 0.75 > \Pr(V^B = i)$ for $i = 0, 100$, that is, asset fundamentals were not independent. In the control treatment (Treatment II), the value of asset B was set equal to the value of C (i.e., $p = 0$), and $\Pr(V^B = i|V^A = i) = 0.5 = \Pr(V^B = i)$, that is, asset fundamentals were independent.

Because in Treatment I asset fundamentals are not independent, informational contagion arises in equilibrium: in the REE, the correlation between prices is 0.71, higher than that between fundamentals (0.5). In contrast, in Treatment II , since the two asset values are independent, informational contagion does not arise in equilibrium—in the REE, the correlation between prices is 0, the same as that between fundamentals.

3.2 Procedures

In each of the 12 sessions of the experiment, we had 10 rounds of trading activity. Each session was organized in the following way:

⁸For the reader's convenience, the correlations are derived in Appendix A.

- Subjects were given written instructions (see Appendix *F*). Subjects could ask clarifying questions, which we answered in private. After reading the instructions, subjects answered a short questionnaire to check their understanding of the experiment. A subject giving the wrong answer was notified that the answer was wrong and was asked to answer again until he gave the correct answer.
- We randomly assigned the 16 participants of each session to two groups of 8 subjects, group *I* and group *II*. A subject remained in the same group for the entire experiment.
- In each round, market *A* and market *B* opened in sequence. In even rounds, group *I* traded asset *A*; after trading in market *A* ended, group *II* traded asset *B*. In odd rounds, group *II* traded asset *A* and group *I* asset *B*.
- In each market, subjects traded the asset by exchanging it among themselves for 200 seconds. They used the trading platform shown in Figure *F1*, Appendix *F*.
- While one group of subjects was trading, the other observed the history of quotes and trades. Moreover, while one of the two groups traded in market *A*, we asked the subjects in the other group to indicate their belief on the value of asset *A* being 100 after 80, 140 and 190 seconds of trading activity. This helped subjects to pay attention to the trading flow in the other market; additionally, it provided us with information about how subjects interpreted the history of trading in market *A*.

Let us now discuss the procedures for each round in detail. In each round, before the start of trading activity in market *A*, the computer program drew the asset value V^A , which with equal probability, was equal to 0 or 100 units of a fictitious experimental currency called lira.

At the beginning of the round each participant received an endowment of 4 units of asset *A* and 500 liras. Subjects also received information about the asset value in the form of a symmetric binary signal with precision 0.75. Specifically, when the asset value was 100, six participants observed a “green ball” and two participants a “red ball;” if the value was equal to 0, six participants observed a “red ball” and two participants a “green ball.” This

signal structure guarantees that in each round the private signals collectively reveal the fundamental V^A even if the number of subjects is finite.⁹

During the 200 seconds of trading activity in market A , subjects could post offers to sell or buy units of asset A . To post a sell offer, a subject would click on a sell button and enter the minimum price he was willing to accept. The offer appeared immediately on everyone's screen, in a column labeled Current Sell Offers (the identity of the subject making the offer was not revealed). Similarly, to post a buy offer, a subject would click on a buy button and enter the maximum price he was willing to pay. A trade would automatically be executed by the computer whenever the lowest sell price (ask) was lower than the highest buy price (bid). As a result, if a subject wanted, for instance, to buy at the prevailing (i.e., the lowest) ask, he could simply enter a price equal to or greater than that price, and the trade would be immediately executed (at the outstanding price). If, instead, a subject input, e.g., an ask price higher than the outstanding ones, his ask would simply appear among the Current Sell Offers (where all asks were shown in increasing order).

Subjects could choose any buy or sell price greater than or equal to zero.¹⁰ For each subject, the maximum number of outstanding sell offers allowed was equal to the units of the asset held in his portfolio; and the sum of all the buy offer prices could not exceed the cash held in his portfolio. At any time, a subject could withdraw outstanding buy or sell offers that had not already been executed.

A subject's screen also displayed his current portfolio of cash and of units of the asset, the list of past trades in the round (with his own executed trades highlighted), all the outstanding bid and ask prices, and the time left before the end of the round (see Figure $F1$).

After trading in market A ended, trading in the other market occurred according to the same protocol. In particular, each participant in market B received an endowment of 4 units of asset B and 500 liras. Subjects also received information about the realization of the random variable C (which we labeled the " B -coin," in the experiment) in the form of a symmetric binary signal with precision 0.75, exactly as explained above for asset A . In the instructions, for Treatment I , we explained to subjects that the value of

⁹Other signal structures, even if informative, may not deliver the same result (for instance, i.i.d. signals with precision 0.75).

¹⁰Also, a subject was not allowed to place a buy offer higher than one of his outstanding sell offers (in other words, a subject could not trade with himself).

asset B was equal either to asset A or to the B -coin with equal probability; and for Treatment *II*, we explained that the value of asset B was equal to the B -coin. At the end of the round, the values of the two assets, V^A and V^B , were revealed, and each subject saw a detailed summary of his per-round payoff on the screen. Each subject's per-round payoff was equal to the sum of the cash and of the value of the assets in his portfolio. Additionally, we paid subjects a transaction bonus of 5 liras for the first 5 trades. The bonus gave subjects an incentive to exchange the asset in an environment in which payoffs and endowments are the same.¹¹ We limited the bonus to the first five trades to avoid the possibility that subjects would keep exchanging the asset among themselves just to earn the bonus.

In each session, before we ran the actual experiment, we had a training phase to familiarize subjects with the trading platform.¹² The training phase consisted of 10 rounds of trading in only one market. Since the session was for training purposes only, we do not report its results in the main text of the paper.¹³

At the end of the experiment, we randomly selected three rounds and summed up the per-round payoffs. We converted experimental liras into British Pounds at the exchange rate of $\pounds 1 = 100$ liras; additionally, subjects earned $\pounds 5$ as a show-up fee. We paid subjects in private immediately after the end of the experiment. On average, subjects earned $\pounds 28$ (approximately equal to $\$47.5$). Sessions lasted approximately 3 hours.

¹¹Note that, in the economy we described in Section 2, there are no gains from trade and as a result, agents do not have any incentive to trade at the REE price (no-trade theorem).

¹²The use of experienced subjects is typical in trading experiments with double auctions, since convergence to competitive equilibrium requires repetitions, even in simple environments (Smith, 1962). For example, Copeland and Friedman (1991) use subjects who had already previously participated in an asset market trading experiment; Forsythe and Lundholm (1990) make subjects participate in double auctions experiments in two consecutive nights. In our experiment, we are interested in the informational spillover from market A to market B ; it is, therefore, important to avoid that the trading activity in market A is only noise and reveals no information about the asset value because subjects are still learning how to trade.

¹³In Appendix *B*, we show that the training session was indeed useful for subjects to learn how to trade.

4 Results

We now turn to the discussion of the experimental results. We will first discuss the results of Treatment *I*, where $p > 0$ and therefore theory predicts that informational contagion occurs. Later, we will compare these results with those of Treatment *II*, where $p = 0$ and, according to theory, there is no informational contagion.

In our model, contagion is caused by an informational spillover from market *A* to market *B*. Of course, the information that subjects in market *B* obtain from the trading activity in market *A* depends on how well the price in that market aggregates private information. For this reason, as a first step, in the next section we investigate the aggregation of private information in market *A*.

4.1 Trading and price convergence in market *A*

As we explained above, given the signal structure implemented in the laboratory, the signals that subjects receive reveal V^A : that is, there is enough information in the market to learn V^A through trading activity. Therefore, the price aggregates private information if, at least by the end of the round, it equals V^A (or, given our transaction bonus, it differs from it by at most 5 liras).

To study the aggregation of information in market *A*, we consider the average of the last five trade prices in each round, which we label the “final price” and denote by \bar{p}_A^{Last5} .¹⁴ Figure 1 shows a histogram of the per-round distance (defined as the absolute value of the difference) between the final price and V^A . As the figure shows, in almost 70 percent of the cases, the distance is less than 20 liras. The distribution is heavily skewed to the right: the average distance is 23 liras, higher than the median, which is less than 10 liras (see Table 1); in 20% of the rounds, the distance is higher than 50 liras. In other words, the price does aggregate private information well, but there are some rounds in which aggregation fails.

¹⁴On average, in each round, there were 20 trades (with a median of 18.5 and a standard deviation of 6.6). Recall that we paid subjects for the first five trades they executed (in order to give them an incentive to trade in an economy in which otherwise there would be no gains from trade). Given that in each market there are 8 subjects and that a trade involves two parties, 20 trades per round is the amount of trading activity that theory suggests our bonus should have generated.




Figure 1: Per-round distance between the final price and V^A . The final price is defined as the average of the last five trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.

These findings are confirmed by regression analysis.¹⁵ As the first column of Table 2 shows, if we regress the final price on V^A , the slope coefficient is positive and significant, but only equal to 0.5: only half the information that the subjects receive is aggregated by the final price.¹⁶ According to the regression, when the fundamental is 0, the expected final price is 23; and when the fundamental is 100, it is 78. In both instances, the final price moves from its unconditional expected value of 50 toward the realization of the fundamental, but it is further away from it than the 5 liras bonus justifies.

Note that our results do not depend on how we define the per-round final price. As the second and third columns of Tables 1 and 2 show, the regression coefficients do not meaningfully change if we define the final price as the average price of the last 3 trades, or the average price of the trades

¹⁵Throughout the paper, for all regression results, we cluster standard errors at the session level and report them in parenthesis. Moreover, we indicate significance at the 10, 5, 1 percent level with *, **, ***, respectively.

¹⁶We reject the hypothesis that the coefficient is equal to the theoretical value of 1—p-value equal to 0.01, using a cluster-robust t-test at the session level.

	Last 5 trades	Last 3 trades	Last 30 seconds of trade
Mean	22.64	22.67	21.67
Median	9.29	8.67	8.60
Std. Dev.	27.86	28.51	28.37
N	60	60	52

Table 1: Distance of the final price of asset A from V^A

The table shows the mean, median and standard deviation of the distance of the final price of asset A from V^A . The final price is the average of the last five trade prices in a round (column 1), or of the last three trade prices (column 2), or of the trades occurred in the last 30 seconds of a round (column 3).

	Last 5 trades	Last 3 trades	Last 30 seconds of trading
V^A	0.545**(0.124)	0.546**(0.126)	0.564**(0.114)
Constant	23.10**(8.380)	22.70**(8.302)	22.56**(8.210)
R-squared	0.475	0.464	0.496
N	60	60	52

Table 2: Regression results for market A

The table shows the regression results of the final price of asset A on V^A . The final price is the average of the last five trade prices in a round (column 1), or of the last three trade prices (column 2), or of the trades that occurred in the last 30 seconds of a round (column 3).

that occurred over the last 30 seconds of trading activity.¹⁷

In summary, our results show that final prices in market A are a noisy signal of the asset value. Therefore, they could be used by subjects in market B to infer V^A and, in turn, to construct their beliefs on V^B . To understand how rational agents formed their beliefs on V^A by observing the prices in market A , we regressed V^A on the final price in a probit regression, and estimated the conditional expected value of V^A (i.e., the conditional probability that $V^A = 100$). Figure 2 shows the conditional expected value of V^A as a function of market A 's final price. For instance, the conditional expected value of V^A is around 28 for a final price of 20 and climbs to over 80 for a

¹⁷See also Appendix C. Note that in eight trading rounds, no transactions occurred in the last 30 seconds of trade. This explains why in column 3 of Tables 1 and 2 the average final price is computed over 52 rather than 60 rounds.




Figure 2: Expected value of V^A as a function of the final price in market A .

final price of 80.¹⁸

Note that the expectations we obtain through the probit regression are very similar to those that a Bayesian agent would compute after observing the empirical frequencies of V^A being equal to 0 or 100 for different ranges of the final price (see the first three columns of Table 3). Interestingly, such empirical expectations are also very close to the belief we elicited from subjects trading in market B while they were observing the trading activity in market A (last column of Table 3).¹⁹

Until now, we have focused on the price towards the end of the 200 seconds of trading. Aggregation of information took some time to occur. In Figure 3, we show the evolution of the distance between the price and V^A over time. We divided the trading round into 10 intervals and computed the average distance in each of these intervals. As the figure shows, the price becomes

¹⁸We report the coefficients of the probit regression in Table C1 in Appendix C.

¹⁹Recall that while one group was trading asset A , subjects in the other group had to state their belief about the value of asset A (i.e., their expected value of the asset) when there was a remaining trading time of 120, 60 and 10 seconds in market A . Here, and in the rest of the paper, we focus on subjects' elicited expectation when there were only 10 seconds of trading activity left in market A . In Appendix D we carry out some additional analysis of subjects' elicited beliefs.

	Frequencies		Bayesian updates	Beliefs
	$V^A = 0$	$V^A = 100$	$\Pr(V^A = 100 \bar{p}_A^{Last5})$	
$\bar{p}_A^{Last5} > 75$	0.09	0.71	0.89	85.8
$50 < \bar{p}_A^{Last5} \leq 75$	0.09	0.08	0.46	59.7
$25 \leq \bar{p}_A^{Last5} \leq 50$	0.09	0.10	0.54	43.0
$\bar{p}_A^{Last5} < 25$	0.73	0.10	0.13	10.7

Table 3: Empirical Bayesian Updates

The table shows: 1) the frequencies with which the final price (defined as the average of the last five trade prices in a round) belonged to a particular range, conditional on V^A (columns 1 and 2); 2) the Bayesian updates about V^A computed using these frequencies (column 3); 3) average subjects' beliefs elicited 10 seconds before the end of the trading activity in market A (column 4).

closer and closer to the fundamental, as private information is aggregated through the trading activity.

4.2 Trading and price convergence in market B

We now turn our attention to market B . We conduct a similar analysis to that of market A , and study the behavior of the final price, defined, as in the previous section, as the average of the last five trade prices. Figure 4 shows the histogram of the per-round distance between the final price in market B and V^B ; Table 4 shows the mean, the median and the standard deviation across rounds. As one can observe, the distance between price and fundamental is higher than in market A (the difference is significant at the 5% level using a Wilcoxon signed-rank test).

This is not surprising. In contrast to market A , the information that subjects receive (i.e., the signals about C and the history of trading activity in market A) does not reveal V^B : there is not enough information in the market to learn the fundamental through trading activity.

The higher distance between final price and fundamental may be due to two reasons: i) the fact that in market B there is less information about the fundamental than in market A ; and ii) the fact that the aggregation of private information in market B is less efficient (e.g., because subjects, when interpreting their private signal, have an additional source of information, the public information from market A , that may confuse them). To gauge




Figure 3: Distance of the price of asset A from V^A over time. The figure shows the distance between the average price and V^A for each interval of 20 seconds.




Figure 4: Per-round distance between the final price and V^B . The final price is defined as the average of the last five trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.

	Last 5 trades	Last 3 trades	Last 30 seconds of trading
Mean	29.93	29.54	29.01
Median	16.7	13.5	15.75
Std. Dev.	29.18	29.68	30.26
N	60	60	58

Table 4: Distance of the final price of asset B from V^B

The table shows the mean, median and standard deviation of the distance of the final price of asset B from V^B . The final price is the average of the last five trade prices in a round (column 1), or of the last three trade prices in a round (column 2), or of the trades occurred in the last 30 seconds of a round (column 3).

the ability of subjects in market B to aggregate information, we compare the final price to the REE price. In the REE, trading activity in markets A and B reveal both V^A and C , but does not reveal whether V^B equals V^A or C . For this reason, as discussed above, the REE price in B is 100 when $V^A = C = 100$; 0 when $V^A = C = 0$; and 50 when $V^A \neq C$. If the (final) price in market B aggregates the information contained in the patterns of trading activity in market A and in the signals, it should equal the REE price.

	Last 5 trades	Last 3 trades	Last 30 seconds of trading
Mean	25.90	25.92	25.32
Median	19.50	19.0	22.33
Std. Dev.	22.30	23.21	23.63
N	60	60	58

Table 5: Distance of the final price of asset B from the REE price

The table shows the mean, median and standard deviation of the distance of the final price of asset B from the REE price. The final price is the average of the last five trade prices in a round (column 1), or of the last three trade prices in a round (column 2), or of the trades occurred in the last 30 seconds of a round (column 3).

As Figure 5 and Table 5 show, the average distance between the final price and the REE price is 25.9. This is not statistically different from the distance in market A (Wilcoxon signed-rank test at session level—p-value = 0.20). In other words, the aggregation of information in market B is not reduced by the fact that the task that subjects are facing is harder.

In order to understand the aggregation of information in market B , we




Figure 5: Per-round distance between the final price and the REE. The final price is defined as the average of the last five trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.

also regress the initial and the final price in market B over the realization of C and over the value of the fundamental in market A . The initial price is defined as the average price of the first 5 trades in each round, whereas the final price is defined, as before, as the average of the last 5.

The results of the initial-price regressions are reported in the first two columns of Table 6. As the first column shows, the effect of C on the initial price is positive and significant; it is also smaller than it would be if signals were immediately reflected in the price at the beginning of the round (0.16 versus 0.5). This is not surprising, as the aggregation of private information happens over time. Additionally, the effect of the fundamental in A (V^A), although positive, is both lower than what theory predicts (0.15 versus 0.5) and non significant. To understand this result, let us look at the second column of Table 6, where we replaced V^A with its conditional expectation $E^{probit}(V^A)$ given by the probit regression (as illustrated in Section 4.1). The coefficient on $E^{probit}(V^A)$ is significant and close to the theoretical value of 0.5. That is, subjects trading in market B correctly incorporate the information coming from market A . It is only because market A 's price is a noisy

	Initial Price (first 5 trades)		Final Price (last 5 trades)	
C	0.165**(0.043)	0.128**(0.038)	0.278**(0.058)	0.235**(0.047)
V^A	0.151(0.083)		0.305*(0.133)	
$E^{probit}(V^A)$		0.418**(0.065)		0.632**(0.137)
Constant	33.49**(6.679)	18.78**(6.202)	12.52(7.825)	-5.616(4.391)
R-squared	0.235	0.463	0.385	0.563

Table 6: Regression results for market B

The table shows the regression results of the initial (final) price of asset B on C and on V^A in column 1 (3). Columns 2 and 4 show the regression results using the conditional expectation of V^A (computed by the probit regression).

signal of V^A that the spillover from market A to market B is lower than the theoretical prediction.

Let us now consider how the final price aggregates subjects' information. In the last two columns of the table, we regress the final price on the realizations of C and on V^A (column 3) or on its conditional expected value (column 4). If the final price aggregated subjects' signals correctly, the coefficient on C should be 0.5 (since $V^B = C$ only with probability 0.5). In both columns, the coefficient is positive and significant, but approximately only half the theoretical value. This is a similar result to what was observed in market A , where the coefficient of 0.5 was half its theoretical counterpart.²⁰ In both markets, the price aggregates subjects' private signals only partially.

Moreover, according to the theory, the coefficients on V^A should be 0.5 (since $V^B = V^A$ with probability 0.5); in the regression, the coefficient is positive and significant, but smaller than the theoretical one. Similarly to what happened for the initial price, however, when we use as a regressor the conditional expectation of V^A the coefficient increases to approach its theoretical value.²¹ In other words, also when we look at the final price, it is apparent that subjects incorporate the information coming from the other market correctly.

Overall, the regression results suggest that the aggregation of public in-

²⁰In market A , full aggregation of private information would have implied that the coefficient on V^A was equal to 1.

²¹We cannot reject the hypothesis that the coefficient is equal to the theoretical one—p-value = 0.38—using a cluster-robust t-test at the session level.

formation from market A occurs in accordance with the theory, both qualitatively and quantitatively. In contrast, both in market A and in market B , subjects have more difficulties aggregating private information. Although private signals are aggregated by the price, the aggregation is not complete.

4.3 Contagion

The previous section clearly documents the existence of an informational spillover from market A to market B . In the literature, financial contagion is usually characterized as “excess correlation” among asset prices: in particular, there is contagion between two markets when the correlation between asset prices is greater than that between asset fundamentals: $Corr(P^A, P^B) > Corr(V^A, V^B)$. In our theoretical model, the informational spillover generates contagion: as we discussed above, the correlation between fundamentals is 0.5, whereas that between prices is 0.71. Contagion also occurs in the laboratory: across rounds, the correlation between final prices is 0.67. This is an important result. It shows that the informational contagion predicted by our theoretical model is also the outcome of subjects’ interactions in a market setting: in other words, the trading strategies subjects put in place in the laboratory lead to a contagion effect from market A to market B .

Not only is the price correlation in the laboratory higher than that between fundamentals, it is also very close to the theoretical one (0.67 is not significantly different from 0.71).²² This is somehow surprising since we know from the previous analysis that the prices observed in the laboratory do not fully aggregate private information (whereas the REE prices do). To shed light on this result, recall two observations that we made in the previous section: i) in the two markets, signals are aggregated only partially, and the level of aggregation is (approximately) similar; ii) subjects attach (approximately) the theoretically correct weight to the information coming from market A . Intuitively, the first observation implies that, holding constant the informational spillover across markets, the variances of P^A and P^B and their covariance are lower than what is predicted by the theory. The second observation implies that the reduction in the covariance (with respect

²²We regressed the final price in market B on the final price in market A multiplied by the ratio of the standard deviations of the two prices (and on a constant). We cannot reject the null hypothesis that the estimated coefficient, which is equal to the correlation index, is equal to 0.71 using a cluster-robust t statistic—p-value equal to 0.8.

to the theoretical one) exactly offsets the reduction in the variances of P^A and P^B , thus leading to a correlation very close to the theoretical one. A simple model in which these two observations hold (with no approximation) is: $P^A = \alpha + \beta V^A + \varepsilon$ and $P^B = \delta + \frac{1}{2}P^A + \frac{1}{2}\beta C + \eta$ (with α and δ being two constants, $0 < \beta < 1$ and ε and η two uncorrelated error terms). It is straightforward to prove that, for $\varepsilon = \eta = 0$, the correlation is identical to that of our theoretical model (in which $\beta = 1$). In the experiment, of course, both error terms have a positive variance, but their net effect on the correlation turns out to be negligible.²³

In the experiment, as in the theoretical model, the information coming from market A increases the market's informational efficiency. Indeed, it is straightforward to show that if subjects in market B had attached zero weight to the information coming from market A (i.e., if the coefficient on $E(V^A)$ in the regression reported in Table 5 had been zero), the distance between the price and the fundamental would have been higher. Nevertheless, in those rounds when the value of asset B differs from that of asset A , the information coming from market A is detrimental: the distance between price and fundamental in the experiment was 40—higher than 27, the average distance when the asset values were equal, and of 29, the average distance across all rounds.²⁴ This shows the negative impact of information contagion: although the information coming from other markets' prices is on average valuable, it becomes counterproductive when price changes reflect idiosyncratic shocks in those markets.

4.4 Independent fundamentals and absence of contagion: the results of Treatment II

Until now, we have shown that allowing subjects to observe the history of trades in another market generates financial contagion in the laboratory. This empirical result agrees with the theoretical predictions: indeed, the correlation we obtain in the laboratory is remarkable close to the equilibrium one. One may wonder, however, whether in the laboratory contagion is really generated by informational spillovers, as in the theoretical model, or rather whether it is a mere artifact, caused by subjects in market B being influenced

²³In Appendix A, we show that the correlation is decreasing in the variance of η but increasing in the variance of ε .

²⁴There are 20 rounds in which the value of asset B differs from that of asset A .

by the trades and prices in market A , independently of their information content.

To tackle this issue, we ran Treatment II , in which we set $p = 0$; that is, V^A and V^B are independently distributed (V^B being equal to C). According to the theory, since the asset fundamentals are independent, there should be no information contagion.

In Table 7, we present the same regressions we had discussed in Table 6 for Treatment I . The differences between the two treatments are striking. First, both when we look at the behavior of the initial and of the final price, the coefficient on the value of asset A and on its probit expectation are now much smaller than in Treatment I , in fact not significantly different from zero. This is in accordance with theory: no information spillover occurs between markets when asset fundamentals are independent. This result suggests that behavioral biases did not cause the informational contagion observed in Treatment I (as would have been the case if, for instance, subjects in market B were affected by the price in market A independently of its information content); on the contrary, subjects use the information coming from market A only when it is relevant.²⁵

	Initial Price (first 5 trades)		Final Price (last 5 trades)	
C	0.291***(0.037)	0.282***(0.036)	0.674***(0.071)	0.670***(0.069)
V^A	0.052*(0.025)		0.013(0.065)	
$E^{probit}(V^A)$		-0.025(0.087)		-0.146*(0.060)
Constant	38.64***(4.621)	42.65**(6.435)	14.43*(6.900)	21.99**(7.102)
R-squared	0.427	0.414	0.627	0.640

Table 7: Regression results for Treatment II

The table shows the regression results of the initial (final) price of asset B on C and on V^A in column 1 (3). Columns 2 and 4 show the regression results using the conditional expectation of V^A (computed by the probit regression).

Additionally, when we look at the final prices, the coefficient on C is now much higher than what was reported for Treatment I (0.67 vs. 0.28). As a matter of fact, a statistical test reveals that this coefficient is not significantly different from the coefficient on the value of asset A in market A (in either

²⁵Subjects' behavior in market A is similar to that of Treatment I . We report some descriptive statistics and regression results in Appendix E .

treatment).²⁶ This happens because the distribution of V^B (equal to C in this treatment) is the same as that of V^A , and subjects disregard the information coming from market A . Similarly to what we observed in market A , in both treatments, private signals are only imperfectly aggregated by the final price (the coefficient of C is less than 1).

Given these results, it is not surprising that the correlation between final prices is -0.13 , not statistically different from that between fundamentals (i.e., zero).²⁷ In other words, when asset values are independent, we do not observe financial contagion in the laboratory.

To gain more intuition on these aggregate results, we now look at how subjects set prices in the two treatments. We regress, separately for the two treatments, the bid and ask prices that subjects posted in market B on their private information and on their elicited belief that $V^A = 100$. The results are reported in Table 8. In Treatment I , the posted bid and ask prices are positively and significantly related not only to a subject's private information, but also to his belief about V^A (see columns 1 and 2). Subjects' beliefs have a relatively large effect: a 10 unit increase in subjects' beliefs about V^A results in, approximately, a 4 unit increase in the bid and in the ask price subjects post in market B . In contrast, in Treatment II , only a subject's private information matters, and not his assessment of the history of market A ; indeed, the coefficients on subjects' beliefs are not significant both for the bid and the ask prices. We obtain analogous results if we look at the probability that a subjects posts a bid rather than an ask as a function of his private information and his belief about V^A through a probit regression: a subject with a high expectation on V^A is significantly more likely to post a bid (that is, to try and buy the asset) in Treatment I ; this effect disappears in Treatment II (see Table 9).²⁸

Taken together, the results of Tables 8 and 9 show that, when posting bid and ask prices, subjects react to the history of trading activity in market A when such history carries information on V^B (as in Treatment I), but disregard it when it does not (as in Treatment II). As a result of their

²⁶In both treatments, that coefficient is 0.55 (see Tables 2 and $E2$). The p-values for the test that the coefficients in columns 3 and 4 are not different from 0.55 are equal to 0.14.

²⁷We cannot reject the null hypothesis that the correlation is equal to 0 (by using the same test discussed in footnote 22)—p-value = 0.13.

²⁸We obtain very similar results (available on request) if we run a logit regression or a linear probability model.

behavior, we observe contagion in Treatment *I* but not in Treatment *II*.

	Treatment <i>I</i>		Treatment <i>II</i>	
	Bid	Ask	Bid	Ask
Good Signal	13.82**(4.835)	10.75**(3.101)	32.60**(5.167)	22.73**(4.163)
Belief	0.425**(0.118)	0.376**(0.080)	-0.0512*(0.023)	-0.076(0.043)
Constant	13.91**(3.16)	29.56***(2.61)	28.57**(5.65)	42.89***(5.01)
R-squared	0.333	0.284	0.192	0.098
<i>N</i>	2,244	2,658	2,347	2906

Table 8: Bid and Ask prices conditional on the subject's private signal and belief

The table shows the results from a regression of ask and bid prices on a subject's belief about the value of good A and on a dummy equal to one if the subject's private signal is good.

	Treatment <i>I</i>	Treatment <i>II</i>
Good Signal	0.345***(0.024)	0.571***(0.071)
Belief	0.00408***(0.001)	-0.00062(0.001)
Constant	-0.494***(0.059)	-0.380***(0.048)
<i>N</i>	4902	5253

Table 9: Probability of posting a bid conditional on the subject's private signal and belief

The table shows the results from a probit regression of a dummy equal to one if the quote posted by a subject is a bid on a dummy equal to one if his private signal is good and on his belief about the value of good A.

5 Conclusions

In actual financial markets, traders often interpret price movements in one market as conveying information about asset fundamental values in other markets. In an influential paper, King and Wadhvani (1990) showed that, in a Rational Expectations Equilibrium, these informational spillovers across

financial markets generate financial contagion, a well-established empirical regularity. We tested the predictions of King and Wadhvani (1990) in the laboratory. Our work supports the predictions of the theory. Although in the laboratory private information is not perfectly aggregated, subjects are able to use the public information coming from the other market correctly. As a result, the correlation between asset prices is very close to the theoretical one. In principle, behavioral biases may lead subjects to under-react to public information coming from another market (and focus, instead, on the information about their own market) or, on the contrary, to overreact to it (for instance, a price decline in another market may cause subjects to be more prone to sell); this, however, does not happen in our experiment. Moreover, in the laboratory, we do not observe contagion when theory suggests we should not, that is, when the history of trades and prices in the other market conveys no relevant information. Overall, our experimental results show that the Rational Expectations Equilibrium performs remarkably well in describing financial contagion and the comovement among asset prices generated by informational spillovers. As a result, in future work with field data, one can study informational contagion with confidence that the Rational Expectations Equilibrium provides a good explanation of asset price comovements.

References

- [1] Allen, F. and Gale D. (2000) Financial Contagion, *Journal of Political Economy*, **108**, 1-33.
- [2] Calvo, G. (1999) Contagion in Emerging Markets: When *Wall Street* is a Carrier, mimeo, University of Maryland.
- [3] Cipriani, M. and Guarino A. (2008) Herd Behavior and Contagion in Financial Markets, *B.E. Journal of Theoretical Economics (Contributions)*, **8**, Article 24.
- [4] Copeland, T. E. and D. Friedman (1991) Partial Revelation of Information in Experimental Asset Markets, *Journal of Finance* **46**, 265-295
- [5] Edwards, S. and R. Rigobon (2002) Currency Crisis and Contagion: An Introduction, *Journal of Development Economics* **69**, 307-313.

- [6] Ehrmann, M., Fratzscher, M. and Rigobon, R. (2011) Stocks, bonds, money markets and exchange rates: measuring international financial transmission, *Journal of Applied Econometrics*, **26**, 948-974.
- [7] Eichengreen, B., A. Rose and C. Wyplosz (1996) Contagious Currency Crisis, *The Scandinavian Journal of Economics* **98**, 463-484.
- [8] Forsythe, R., and R. Lundholm. (1990) Information aggregation in an experimental market, *Econometrica* **58**, 309-347.
- [9] Fostel, A. and J. Geanakoplos (2008) Leverage Cycles and The Anxious Economy, *American Economic Review*, **98**, 1211-1244.
- [10] Grossman, S. J. and Stiglitz, J. (1980) On the Impossibility of Informationally Efficient Markets, *American Economic Review*, **70**, 393-408.
- [11] King, M. and Wadhvani, S. (1990) Transmission of Volatility between Stock Markets, *Review of Financial Studies*, **3**, 5-33.
- [12] Kodres, L. and Pritsker, M. (2002) A Rational Expectations Model of Financial Contagion, *Journal of Finance*, **57**, 769-799.
- [13] Kyle, A. and Xiong, W. (2001) Contagion as a Wealth Effect, *Journal of Finance*, **56**, 1401-1440.
- [14] Pavlova, A. and Rigobon, R. (2007) The Role of Portfolio Constraints in the International Propagation of Shocks, *Review of Economic Studies*, **75**, 1215-1256.
- [15] Smith, V. (1962) An Experimental Study of Competitive Market Behavior," *Journal of Political Economy* **70**, 111-137.
- [16] Trevino, I. (2013) Channels of Financial Contagion: Theory and Experiments, Mimeo, New York University.
- [17] Yuan, K. (2005) Asymmetric Price Movements and Borrowing Constraints: A REE Model of Crisis, Contagion, and Confusion, *Journal of Finance*, **60**, 379-411.

Appendix (for online publication)

Appendix A: REE and correlations

REE derivation

In this appendix we derive the REE. To this aim, let us first define the Private Information Equilibrium (PIE), that is, the equilibrium in which each agent only uses his private information and neglects the information contained in the price. Figure A1 shows the PIE in market A when $V^A = 0$. Since the precision of the private signal is 0.75 and 8 agents trade in market A , 6 agents evaluate the asset 25 liras and 2 agents evaluate it 75 liras. Bearing in mind that each agent is endowed with 4 units of the asset and 500 liras, supply and demand curves are easily derived. For instance, let us consider the supply curve. At a price lower than 25, no agent is willing to supply the asset. At a price of 25, 6 agents are just indifferent between holding and selling the asset (the maximum supply is, therefore, 24 units). At any price between 25 and 75, these 6 agents supply all their endowment. At a price of 75, also the other 2 agents become weakly willing to supply the asset. For a price higher than 75, all 32 units are supplied. The PIE price is 41.7, where demand and supply cross. A similar analysis shows that the PIE price when $V^A = 100$ is 75 (as illustrated in Figure A2). Of course, these two prices cannot be a REE. Indeed from the first price, agents infer that the value is 0 and from the second, that it is 100. Therefore in the first case demand and supply become those illustrated in Figure A3; and in the second case they look like in Figure A4. The REE prices are 0 and 100 respectively.

The analysis for market B follows the same logic. In Figure A5 we shows the REE (assuming $p = 0.5$) when $V^A \neq C$. The equilibrium price becomes 50.




Figure A1: Private Information Equilibrium conditional on $V^A = 0$




Figure A2: Private Information Equilibrium conditional on $V^A = 100$




Figure A3: Rational Expectations Equilibrium conditional on $V^A = 0$




Figure A4: Rational Expectations Equilibrium conditional on $V^A = 100$




Figure A5: Rational Expectations Equilibrium conditional on $V^A \neq C$.

Correlations

We now turn to the computation of the correlation coefficients. We start with those presented in Section 2.2. First, let us derive the correlation between the fundamentals. To do so, we compute the variances and covariance of V^A and V^B :

$$Var(V^A) = E(V^{A^2}) - [E(V^A)]^2 = \frac{1}{2}100^2 - 50^2 = 2500$$

$$\begin{aligned} Var(V^B) &= E(V^{B^2}) - [E(V^B)]^2 = \\ E(V^{B^2}|V^B = V^A) \Pr(V^B = V^A) &+ E(V^{B^2}|V^B = C) \Pr(V^B = C) - [E(V^B)]^2 = \\ \frac{1}{2}100^2 p + \frac{1}{2}100^2(1-p) - 50^2 &= 2500 = Var(V^A) \end{aligned}$$

$$\begin{aligned}
Cov(V^A, V^B) &= E(V^A V^B) - E(V^A)E(V^B) = \\
E(V^A V^B | V^B = V^A) \Pr(V^B = V^A) &+ E(V^A V^B | V^B = C) \Pr(V^B = C) - E(V^A)E(V^B) = \\
\frac{1}{2}100^2 p + \frac{1}{4}100^2(1-p) - 50^2 &= 2500p = pVar(V^A)
\end{aligned}$$

Therefore, the correlation coefficient between fundamentals is:

$$Corr(V^A, V^B) = \frac{pVar(V^A)}{\sqrt{Var(V^A)}\sqrt{Var(V^A)}} = p.$$

We now turn to the computation of the correlation coefficient between prices. In the REE, $P^A = V^A$ and $P^B = pV^A + (1-p)C$. Therefore, variances and covariance are:

$$Var(P^A) = Var(V^A) = 2500$$

$$\begin{aligned}
Var(P^B) &= Var(pV^A + (1-p)C) = \\
p^2Var(V^A) + (1-p)^2Var(C) + 2Cov(V^A, C) &= Var(V^A)(1-2p+2p^2)
\end{aligned}$$

$$\begin{aligned}
Cov(P^A, P^B) &= Cov(V^A, pV^A + (1-p)C) = \\
pVar(V^A) + (1-p)Cov(V^A, C) &= pVar(V^A)
\end{aligned}$$

The correlation coefficient between prices is therefore given by

$$Corr(P^A, P^B) = \frac{pVar(V^A)}{\sqrt{Var(V^A)}\sqrt{Var(V^A)(1-2p+2p^2)}} = \frac{p}{\sqrt{(1-2p+2p^2)}}.$$

It is easy to verify that $\frac{p}{\sqrt{(1-2p+2p^2)}} > p$ for all $0 < p < 1$.

For $p = 0.5$, $Corr(P^A, P^B) = 0.71$.

Now we turn to the computation of the correlation coefficient for the model presented in Section 4.3.

Consider the case in which the prices in both markets aggregates the information only partially. In particular, suppose that

$$\begin{aligned} P^A &= \alpha + \beta V^A \text{ and} \\ P^B &= \delta + pP^A + (1-p)\beta C. \end{aligned}$$

In this case, the variances of the prices and their covariance can be expressed as follows:

$$\text{Var}(P^A) = \beta^2 \text{Var}(V^A).$$

$$\text{Var}(P^B) = \beta^2 \text{Var}(pV^A + (1-p)C) = \beta^2(1-2p+2p^2)\text{Var}(V^A).$$

$$\begin{aligned} \text{Cov}(P^A, P^B) &= \text{Cov}(\beta V^A, \beta(pV^A + (1-p)C)) = \\ &= \beta^2 [p\text{Var}(V^A) + (1-p)\text{Cov}(V^A, C)] = \beta^2 p\text{Var}(V^A). \end{aligned}$$

Therefore, the correlation coefficient is

$$\text{Corr}(P^A, P^B) = \frac{\beta^2 p \text{Var}(V^A)}{\sqrt{\beta^2 \text{Var}(V^A)} \sqrt{\beta^2(1-2p+2p^2)\text{Var}(V^A)}} = \frac{p}{\sqrt{(1-2p+2p^2)}},$$

which is equal to the correlation obtained above for the REE. The partial information aggregation does not affect the correlation since it affects variances and covariance in the same way.

Finally, we compute the correlation coefficient when

$$P^A = \alpha + \beta V^A + \varepsilon$$

and

$$P^B = \delta + \frac{1}{2}P^A + \frac{1}{2}\beta C + \eta$$

(note that, consistently with the text, we are considering the case in which $p = \frac{1}{2}$).

In this case, the variances of the prices and their covariance can be expressed as follows:

$$\text{Var}(P^A) = \beta^2 \text{Var}(V^A) + \text{Var}(\varepsilon).$$

$$\text{Var}(P^B) = \frac{1}{4}\beta^2 \text{Var}(V^A) + \frac{1}{4}\text{Var}(\varepsilon) + \frac{1}{4}\beta^2 \text{Var}(C) + \text{Var}(\eta)$$

$$Cov(P^A, P^B) = \frac{1}{2}\beta^2 Var(V^A) + \frac{1}{2}Var(\varepsilon).$$

Therefore, the correlation coefficient is

$$Corr(P^A, P^B) = \frac{\beta^2 Var(V^A) + Var(\varepsilon)}{\sqrt{((\beta^2 Var(V^A) + Var(\varepsilon)) (2\beta^2 Var(V^A) + Var(\varepsilon) + 4Var(\eta))}}.$$

Finally note that this expression is decreasing in $Var(\eta)$ but increasing in $Var(\varepsilon)$.

Appendix B: Results for the training phase of the experiment

Recall that in each session, and for both treatments, there was a training phase with the purpose of familiarizing subjects with the trading platform. The training phase consisted of 10 rounds of trading in one market; the trading protocol was identical to the one we used for market A in the real experiment.

In this appendix, we show that the training phase was useful for subjects to familiarize themselves with the trading platform and to learn how to trade in a market. To this aim, we compare how the private information was aggregated in the first and in the last five rounds of the training phase. Since the training phase was identical in Treatment *I* and *II*, we pooled together all rounds from both treatments. Figures *B1* and *B2* show the distance of the final price from the asset fundamental value in the first and in the last five rounds, respectively. As can be easily observed, the price aggregates private information to a greater extent in the last five rounds. A Kolmogorov-Smirnov test confirms this result: the hypothesis that the two distributions in Figures *B1* and *B2* are the same is rejected (p-value equal to 0.001).




Figure B1: Distance between the final price and V^A in the first five rounds of the training phase. The final price is defined as the average of the last five trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.




Figure B2: Distance between the final price and V^A in the last five rounds of the training phase. The final price is defined as the average of the last five trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.

Appendix C: Additional results and robustness checks

In this section we report additional results on the probit regression discussed in Section 4.1, and some robustness checks.

Probit regression of V^A on the final price

	V^A
\bar{p}_A^{Last5}	0.029***(0.007)
Constant	-1.164***(0.294)
Pseudo R -squared	0.406
N	60

Table C1: Probit regression for market A

The table shows the results of a probit regression of asset A' s value on the final price.

Robustness checks for market A

This subsection shows the histograms of the distance between the final price of asset A and the fundamental value using different definitions for the final price. In Figure $C1$ the final price is computed as the average of the last three trade prices in a round. In Figure $C2$ it is computed as the average of the trade prices in the last 30 seconds of trade in a round.




Figure C1: Per-round distance between the final price and V^A . The final price is defined as the average of the last three trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.




Figure C2: Per-round distance between the final price and V^A . The final price is defined as the average trade prices in the last 30 seconds of trade in a round. The mean is indicated by the solid line; the median by the dashed line.

Robustness checks for market B

This subsection shows the histograms for market B of the distance between the final price and V^B (Figures $C3 - C4$) and of the distance between the final price and the REE price (Figures $C5 - C6$), using different definitions for the final price. In Figure $C3$ and $C5$ the final price is defined as the average of the last three trade prices in a round, whereas in Figure $C4$ and $C6$ it is defined as the average of the trade prices in the last 30 seconds of trade in a round.




Figure C3: Per-round difference between the final price and V^B . The final price is defined as the average of the last three trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.




Figure C4: Per-round difference between the final price and V^B . The final price is defined as the average trade prices in the last 30 seconds of trade in a round. The mean is indicated by the solid line; the median by the dashed line.




Figure C5: Per-round distance between the final price and the REE. The final price is defined as the average of the last three trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.




Figure C6: Per-round distance between the final price and the REE. The final price is defined as the average trade prices in the last 30 seconds of trade in a round. The mean is indicated by the solid line; the median by the dashed line.

Appendix D: Analysis of beliefs

Recall that in our experiment, while a group of subjects traded in market A , the other group (who would later trade in market B) observed market A prices and trading activity. While they were doing so, subjects in the latter group were asked to report their belief on the value of asset A being 100. They had to do so in three occasions: when the remaining trading time was 120, 60 and 10 seconds. In this appendix, we present a brief analysis of these data. For expositional convenience, we will sometimes refer to the beliefs when the remaining trading time was 120, 60 and 10 seconds as the initial, intermediate and final beliefs.

First of all, it is instructive to look at the evolution of beliefs over time. Figures $D1$ and $D2$ show the average distance (defined as the absolute value of the difference) between asset A 's value and the subjects' beliefs respectively for Treatment I and II . As one would expect, subjects' expectations tend to approach the value of the fundamental as time goes by. The distance of the value of asset A from the initial belief is around 33, whereas that from the final belief is only 25. As we know, the price in market A converged over time to the value of the asset; as this happened, also subjects in group B made better predictions on V^A . Nevertheless, since price aggregation was not perfect, subjects' beliefs at the end of the round are still 25 units far from V^A .

To understand better how subjects form their expectations, we computed the difference between the subjects' final belief and the average of all trade prices occurring between the intermediate and the final belief. The final belief follows the price observed in market A : in 62% of the cases, the difference between the belief and the price was between -10 and 10 units. In Table $D1$, we report the results of a regression of subjects' beliefs on asset A average price in the 50 seconds before the belief elicitation. The upper (lower) panel refers to Treatment I (Treatment II). The coefficients are all statistically significant and vary between 0.66 and 0.87, indicating that beliefs followed the observed prices quite closely.

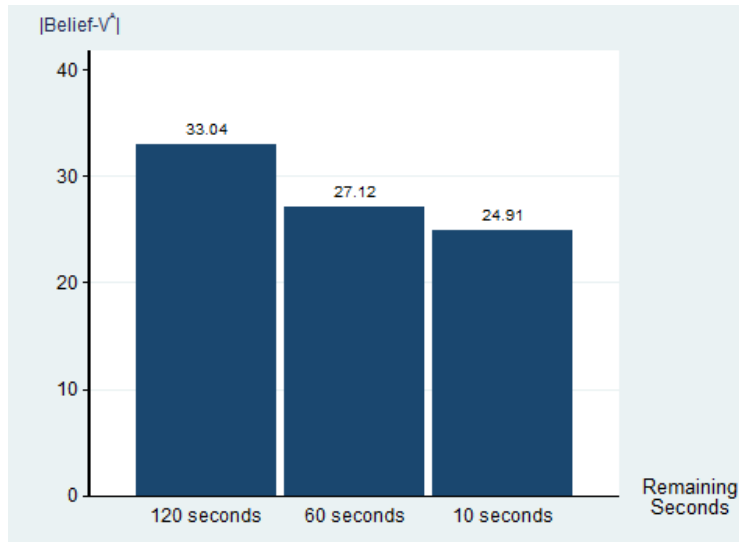


Figure D1: Distance between subjects' beliefs and V^A in Treatment I.




Figure D2: Distance between subjects' beliefs and V^A in Treatment II.

	Initial belief	Intermediate belief	Final belief
	Treatment <i>I</i>		
Price	0.748***(0.101)	0.837***(0.064)	0.869***(0.044)
Constant	10.96(8.499)	6.171(4.626)	4.289(3.341)
R-squared	0.57	0.73	0.74
<i>N</i>	464	464	463
	Treatment <i>II</i>		
Price	0.656***(0.035)	0.774***(0.025)	0.831***(0.019)
Constant	16.71***(2.198)	10.40**(2.657)	8.825**(1.672)
R-squared	0.49	0.64	0.68
<i>N</i>	478	456	472

Table D1: Regression of subjects' beliefs on trade prices

The table shows the regression results of subjects' initial (column 1), intermediate (column 2) and final (column 3) beliefs on market A' s prices. The upper (lower) panel shows the results for Treatment *I* (Treatment *II*).

Appendix E: More results for Treatment II

In the text, when we described Treatment *II*, we only reported some results for market *B*, which is the main object of our interest. In this appendix we report the results for market *A* and some additional results for market *B*.

Distance between the final price and the asset value in market A

Table *E1* reports the mean, median and standard deviation for the distance between the asset value and the final price (defined in three different ways). Figure *E1* reports the histogram of this distance (when the final price is computed as the average of the last five trade prices).

	Last 5 trades	Last 3 trades	Last 30 seconds of trading
Mean	22.16	21.34	22.18
Median	5.6	4.5	4.25
Std. Dev.	30.40	30.77	33.03
<i>N</i>	60	60	56

Table E1: Distance of the final price of asset A from V^A

The table shows the mean, median and standard deviation of the distance of the final price of asset A from V^A . The final price is the average of the last five trade prices in a round (column 1), or of the last three trade prices (column 2), or of the trades occurred in the last 30 seconds of a round (column 3).




Figure E1: Per-round distance between the final price and V^A . The final price is defined as the average of the last five trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.

Regression results for market A

	Last 5 trades	Last 3 trades	Last 30 seconds of trading
V^A	0.555**(0.101)	0.572**(0.092)	0.556**(0.092)
Constant	21.06*(8.296)	20.33**(7.898)	21.93**(8.054)
R-squared	0.458	0.467	0.419
N	60	60	56

Table E2: Regression results for market A

The table shows the regression results of the final price of asset A on V^A . The final price is the average of the last five trade prices in a round (column 1), or of the last three trade prices (column 2), or of the trades occurred in the last 30 seconds of a round (column 3).

Empirical Bayesian updates

Table E3 reports the frequencies of cases in which the last price was in a specific interval. It also reports the beliefs of a Bayesian agent relying on

these frequencies. Figure E2 shows the conditional expected value of V^A obtained from a probit regression of V^A on the final price.

	Frequencies		Bayesian updates	Beliefs
	$V^A = 0$	$V^A = 100$	$\Pr(V^A = 100 \bar{p}_A^{Last5})$	
$\bar{p}_A^{Last5} > 75$	0.09	0.68	0.88	89.0
$50 < \bar{p}_A^{Last5} \leq 75$	0.09	0.18	0.66	63.2
$25 \leq \bar{p}_A^{Last5} \leq 50$	0.06	0.04	0.36	35.0
$\bar{p}_A^{Last5} < 25$	0.75	0.10	0.12	13.9

Table E3: Empirical Bayesian Updates

The table shows: 1) the frequencies with which the final price (defined as the average of the last five trade prices in a round) belonged to a particular range, conditional on V^A (columns 1 and 2); 2) the Bayesian updates about V^A computed using these frequencies (column 3); 3) average subjects' beliefs elicited 10 seconds before the end of the trading activity in market A (column 4).




Figure E2: Expected value of V^A as a function of the final price in market A.

Distance between the final price and the asset value in market B

Table $E4$ reports the mean, median and standard deviation for the distance between the asset value and the final price (defined in three different ways). Figure $E3$ reports the histogram of this distance (when the final price is computed as the average of the last five trade prices).

	Last 5 trades	Last 3 trades	Last 30 seconds of trade
Mean	16.39	15.82	15.03
Median	4.1	3.7	3.7
Std. Dev.	25.97	26.02	24.95
N	60	60	54

Table E4: Distance of the final price of asset A from V^B

The table shows the mean, median and standard deviation of the distance of the final price of asset B from V^B . The final price is the average of the last five trade prices in a round (column 1); of the last three trade prices in a round (column 2); or of the trades occurred in the last 30 seconds of a round (column 3).




Figure E3: Per-round distance between the final price and V^B . The final price is defined as the average of the last five trade prices in a round. The mean is indicated by the solid line; the median by the dashed line.

Appendix F: Instructions

Instructions for the Experiment: Phase I

Welcome to our experiment! You are about to take part in a study on decision making with 15 other participants. The experiment consists of two phases. You will now read the instructions for Phase I and participate in it. For Phase II you will later receive additional instructions.

Everyone has the same instructions. Whenever you have questions, please, do not hesitate to ask one of the supervisors for clarification. Please, do not ask your questions loudly or try to communicate with other participants.

Before the experiment starts, we will randomly assign each of you to one of two groups: half of you (8 participants) will belong to group I, and the other half to group II. You belong to the same group throughout the entire experiment (your group will be shown on the computer screen).

The Experiment

The first phase of the experiment consists of 10 rounds. In every round, participants in each group have the opportunity to trade a good among themselves. Trading lasts for 200 seconds. Participants in each group only observe the decisions made in their group and can only trade among themselves.

The value of the good is expressed in a fictitious currency called “lira,” which will be converted into British Pounds at the end of the experiment according to the following exchange rate:

$$100 \text{ liras} = \text{£}1.$$

The value of the good

At the beginning of every round, the value of the good will be determined by the computer with a mechanism simulating the tossing of a fair coin. The good can have value 0 or 100 liras depending on whether the coin lands heads or tails. Like in the toss of a fair coin, the chances of the good having value 0 or 100 are equal. Note that for each group the computer simulates the tossing of a fair coin at the beginning of every round. Thus, in each round the value of the good is the same for all participants in the same group. The value of the good can, however, change from round to round. And whether the value in a round is 0 or 100 does not depend on the value in previous rounds.

The information you will receive

All participants will receive some information about the value of the good.

How is this information given?

Suppose the value of the good in one group is 100 liras. In this case, we will use a random device similar to an urn with 8 coloured balls: 6 balls are green and 2 are red. Each of the 8 participants will receive one of these balls. Therefore, there is a chance of $3/4$ (equal to $6/8$) that you will observe the message “The colour of the ball is GREEN” on your computer screen; and there is a chance of $1/4$ (equal to $2/8$) that you will observe the message “The colour of the ball is RED”.

Suppose, instead, that the value of the good is 0 liras. In this case, we will use a random device similar to an urn with 8 coloured balls: 6 balls are red and 2 are green. Each of the 8 participants will receive one of these balls. Therefore, there is a chance of $3/4$ (equal to $6/8$) that you will observe the message “The colour of the ball is RED” on your computer screen; and there is a chance of $1/4$ (equal to $2/8$) that you will observe the message “The colour of the ball is GREEN”.

To recap:

- If the value is 100, then there are more GREEN balls in the box.
- If the value is 0, then there are more RED balls in the box.

Therefore, the colour of the ball will give you some information about the value of the good.

When is the information given?

Every participant receives his/her information at the beginning of the 200 seconds.

Procedure for each round

The sequence of activities in each round will be the following:

1. Participants receive information on the value of the good in their group.
2. Participants trade the good for 200 seconds.
3. At the end of the 200 seconds, all participants receive information on the outcomes of their trading activity. In particular, everyone observes the true value of the good and will be able to compute his/her own payoff according to the rules indicated below.

After the first round is concluded, we start the second round of the experiment. The procedures are identical to those of the first round.

Trading

In Figure 1 you see a screen-shot of the trading platform on your computer. In the upper part of the screen, there are two boxes showing the existing Buy Offers and Sell Offers. In the lower part, there are buttons that you can use to buy or sell, and a box where you can insert the price at which you are willing to buy or sell.

On the top left-hand side you can see your holdings of cash and units of the asset (i.e., your Portfolio). On the bottom, you see a continuously updated history of the prices at which the good is traded.

Initial Endowment

At the beginning of each round, you receive an endowment of 4 units of the good and 500 liras. You can use your endowment to trade during the round. The box “Portfolio” is updated whenever you buy or sell a unit of the good. When you buy one unit of the good, the number of units of the good in your portfolio increases by one, and the amount of liras decreases by the price you have paid. When you sell one unit, the number of units of the good in your portfolio decreases by one, and the amount of liras increases by the price at which you have sold.

How to sell or buy

Buying and selling is very simple. If you want to sell one unit of the good, you simply click on the button SELL and enter the minimum amount of liras you want to obtain. Your offer appears immediately in the column Sell Offers where all open sell offers are collected. The open sell offers are ordered with the lowest price being on the top of the list. You can easily identify your own sell offers because they are marked with a button that gives you the opportunity to cancel them, if you so wish.

Similarly, if you want to buy one unit of the good, click on the button BUY and enter the maximum amount of liras you are willing to pay. Your offer appears immediately in the column Buy Offers, where all open buy offers are collected. The open buy offers are ordered with the highest price being on the top of the list. You can easily identify your open buy offers, because they are marked with a button that gives you the opportunity to cancel them, if you so wish.

You are always allowed to withdraw your buy or sell offer that have not been executed: just click on Cancel on the order you want to withdraw.

When and how does a trade take place? A trade is possible if the lowest Sell Price is lower than the highest Buy Price. In this situation, one participant is willing to pay more for the good than another participant asks for it. This situation is recognized by the system and trading takes place automatically.

A simple example will clarify this. Suppose that in a particular moment the lowest Sell Price is 55 liras and the highest Buy Price is 53 liras. Then, no trade is possible. If another participant is willing to buy at 55 liras, the only thing s/he needs to do is enter a Buy Price of 55 liras into the system. The system recognizes that a trade is possible and trade takes place: that is, the seller receives 55 liras from the buyer and the buyer receives one unit of the good from the seller. Note that the transaction always occurs at the pre-existing price. For instance, even if a participant enters a Buy Price of 61 in the system, since the pre-existing lowest Sell Price is 55, the transaction will occur at 55 liras. In other words, if you see a Sell Price at which you are willing to buy, it is enough that you enter a Buy Price equal or greater than that to buy the good.

Consider another example. Suppose that in a particular moment the highest Buy Price is 30 liras and the lowest Sell Price is 37 liras. Then, no trade is possible. If another participant is willing to sell at 30 liras, the only thing s/he needs to do is enter a Sell Price of 30 liras into the system. The system recognizes that a trade is possible and trade takes place: that is, the seller receives 30 liras from the buyer and the buyer receives one unit of the good from the seller. As we said, the transaction always occurs at the pre-existing price. Therefore, even if a participant enters a Sell Price of 23 in the system, since the pre-existing highest Buy Price is 30, the transaction will occur at 30 liras.

As we said, the list of all prices at which a transaction took place appears on the bottom of the screen. The most recent transaction prices are on the top of the list. Your own transactions are identified so that you can keep track of your previous decisions.

Payoff in each round

At the end of every round, you will be told the true value of the good. Your total per-round payoff depends on: 1) the final value of your portfolio, which depends on the value of the good and the amount of liras and the number of units of the good that you hold at the end of the round, and 2) an extra payoff, which depends on the number of trades (sell or buy) you have made during the round.

Value of your portfolio

The value of your portfolio is computed in the following way:

$$\text{Value of portfolio} = \text{liras} + (\text{units of the good}) \times (\text{value of the good})$$

Example 1: Suppose you end a round having 200 liras of cash and 8 units of the good. Suppose the value of the good in that round is 100. Then the value of your portfolio is $200 + (8) * (100) = 1000$.

Suppose, instead, that the value of the good is 0. Then the value of your portfolio is $200 + (8) * (0) = 200$.

Extra payoff

You receive an extra payment of 5 liras for the first 5 buy or sell trades that you execute (i.e., up to a maximum of 25 liras in each round).

Example 1: If in one round you sell 2 goods and buy 2 goods (4 trades), you will earn an additional payment of $4 * 5 = 20$ liras.

Example 2: If in one round you sell 6 goods and buy 1 good (7 trades), you will earn $5 * 5 = 25$ liras, as your extra payment cannot exceed 25 liras.

Note that your extra payment will not immediately increase the amount of liras in your portfolio (which you can use to buy more assets), but will be only part of your final pay-off.

Your total per-round payoff will therefore be:

$$\text{Total per-round payoff} = \text{value of your portfolio} + \text{extra payoff}$$

Payment

This first part of the experiment (Phase *I*) is meant as training for Phase *II*. It gives you the opportunity to learn how to trade. Although we will compute the payoffs as described above, they will not affect your final payment. It is, however, important that you do your best to make profits also in this first part, since what you learn here will be useful for Phase *II*, **which will be paid**. In Phase *II*, the payoffs will be computed in the same way as described above. Those payoffs will be relevant for your final payment: the more money you make by trading, the higher your payment will be. We will convert your liras into pounds at the exchange rate of 100 liras = £1. That is, for every 100 liras you earn in Phase *II*, you will get 1 pound. Moreover, you will receive a participation fee of £5 just for showing up on time. You will be paid in cash (in private) at the end of the experiment.

You will now go through a short questionnaire to make sure that you have understood the instructions and then the experiment will start.

Instructions for Phase II - Treatment I

Let us now move the Phase *II* of the experiment.

Phase *II*

This phase consists of 10 rounds. The rules are identical for all rounds.


As we said, you belong to the same group as in Phase *I*. The main difference with respect to Phase *I* is that there are now two goods, good *A* and good *B*. Moreover, in each round, the two goods are traded one after the other: first, one group trades good *A*; then, after the group has finished, the other group trades good *B*. Each group trades for 200 seconds. Whenever a group is not trading, every participant in that group can observe the trading activity of the other group.

The value of good *A*

As in Phase *I*, at the beginning of every round, the value of good *A* will be determined by the computer with a mechanism simulating the tossing of a coin. The coin can have value 0 liras or value 100 liras depending on whether it lands heads or tails. Like in the toss of a fair coin, the chances of the coin having value 0 or 100 are equal.

The value of good *B*

The value of good *B* will be either 0 or 100 liras. In particular, it will be equal **either** to the value of good *A* (with 50% chance), **or** to the value of a second coin, the “*B*-coin” (also with 50% chance).



The value of the *B*-coin will also be determined by the computer at the beginning of every round, by simulating the tossing of a coin. The coin can

have value 0 liras or value 100 liras depending on whether it lands heads or tails. Like in the toss of a fair coin, the chances of the coin having value 0 or 100 are equal.

In other words, suppose good A is worth 100. Then,

Value of good $B = 100$	with 50% chance
Value of good $B = \text{Value of } B\text{-coin}$	with 50% chance

Suppose, instead that good A is worth 0. Then,

Value of good $B = 0$	with 50% chance
Value of good $B = \text{Value } B\text{-coin}$	with 50% chance

Information you will receive

When good A is traded, all participants belonging to the group trading good A will receive some information about the value of good A . When good B is traded, all participants belonging to the group trading good B will receive some information on the value of the B -coin. However, participants belonging to the group trading good B will not know whether the value of good B is equal to that of good A or whether it is determined by the value of the B -coin.

How is this information given?

When you trade good A , you will receive information on the value of the good exactly as described in Phase I . If the value is 100, you will receive a coloured ball; 6 participants will receive a green ball, whereas only 2 will receive a red ball. If, instead, the value is 0, 6 participants will receive a red ball, whereas only 2 will receive a green ball.

When you trade good B , you will not receive information on the value of the good, but on the value of the B -coin. The procedure will be the same. If the value of the B -coin is 100 you will receive a coloured ball; 6 participants will receive a green ball, whereas only 2 will receive a red ball. If, instead, the value of the B -coin is 0, 6 participants will receive a red ball, whereas only 2 will receive a green ball.

When is the information given?

As in Phase I , every participant receives his/her information before his/her group starts trading (at the beginning of the 200 seconds).

Procedures for each round

As indicated above, the groups *I* and *II* trade in sequence. In odd rounds (1-3-5-...), group *I* trades good *A* (for 200 seconds), and then group *II* trades good *B* (for 200 seconds). In even rounds (2-4-6-...), group *II* trades good *A* (for 200 seconds), and then group *I* trades good *B* (for 200 seconds).

The sequence of activities in round 1 will be the following:

1. Group *I* participants receive information on the value of good *A*.
2. Group *I* participants trade good *A* for 200 seconds, while participants in the other group (*II*) only observe. While observing the behavior of Group *I* participants, Group *II* participants will indicate, on a separate form, what they think the chance is that the true value of good *A* is 100.
3. Group *II* participants receive information on the value of the *B*-coin.
4. Group *II* participants trade good *B* for 200 seconds, while participants of the other group only observe.
5. All participants receive information on the outcomes of their trading activity. In particular, everyone will observe the true value of good *A* and of good *B* and will be able to compute his/her own payoff according to the rules indicated below.

After the first round is concluded, we start the second round of the experiment. The procedures are identical to those of the first round, with the exception that now group *II* starts and trades good *A* and then group *I* trades good *B*. The experiment continues until the 10th round is completed.

Trading

The trading platform on your computer, the initial endowment, and the way you sell or buy a good are all the same as in Phase *I*. The only difference is that now the two groups trade different goods (with possibly different values) one after the other, with one group observing the trading activity of the other.

Payment at the end of the experiment

The per-round payoffs will be determined in the same way as in Phase *I*. At the end of every round, you will be told the true value of both goods *A* and *B*. Your total per-round payoff depends on:

- 1) what you hold at the end of the round: the amount of liras plus the value of the units of the good A or B (according to the good that you traded);
- 2) an extra payment of 5 liras for the first 5 buy or sell trades that you execute (up to a maximum of 25 liras in each round).

We will randomly select 3 out of the 10 rounds of Phase II and we will sum your per-round payoffs in these three rounds to determine your final payoff in liras for Phase II . We will then convert liras into pounds at the exchange rate of 100 liras = £1 and we will sum up this amount to the participation fee of £5. We will pay you in private, immediately at the end of the experiment.

You will now go through a short questionnaire to make sure that you have understood the instructions and then Phase II will start.

Instructions for Phase II - Treatment II

Let us now move the Phase *II* of the experiment.

Phase *II*

This phase consists of 10 rounds. The rules are identical for all rounds.

As we said, you belong to the same group as in Phase *I*. The main difference with respect to Phase *I* is that there are now two goods, good *A* and good *B*. Moreover, in each round, the two goods are traded one after the other: first, one group trades good *A*; then, after the group has finished, the other group trades good *B*. Each group trades for 200 seconds. Whenever a group is not trading, every participant in that group can observe the trading activity of the other group.

The value of good *A*

As in Phase *I*, at the beginning of every round, the value of good *A* will be determined by the computer with a mechanism simulating the tossing of a coin. The coin can have value 0 liras or value 100 liras depending on whether it lands heads or tails. Like in the toss of a fair coin, the chances of the coin having value 0 or 100 are equal.

The value of good *B*

At the beginning of every round, the value of good *B* will be determined by the computer with a mechanism simulating the tossing of a coin. The coin can have value 0 liras or value 100 liras depending on whether it lands heads or tails. Like in the toss of a fair coin, the chances of the coin having value 0 or 100 are equal. Note that the computer will use one coin (coin *A*) for good *A* and one coin (coin *B*) for good *B*. These two coin tosses are independent, that is, the outcome of the one coin toss does not affect the other.

Information you will receive

When good *A* is traded, all participants belonging to the group trading good *A* will receive some information about the value of good *A*. When good *B* is traded, all participants belonging to the group trading good *B* will receive some information on the value of good *B*.

How is this information given?

When you trade good *A*, you will receive information on the value of the good exactly as described in Phase *I*. If the value is 100, you will receive a coloured ball; 6 participants will receive a green ball, whereas only 2 will receive a red ball. If, instead, the value is 0, 6 participants will receive a red ball, whereas only 2 will receive a green ball.

When you trade good B , you will receive information on the value of good B . The procedure will be the same. If the value of good B is 100 you will receive a coloured ball; 6 participants will receive a green ball, whereas only 2 will receive a red ball. If, instead, the value of good B is 0, 6 participants will receive a red ball, whereas only 2 will receive a green ball.

When is the information given?

As in Phase I , every participant receives his/her information before his/her group starts trading (at the beginning of the 200 seconds).

Procedures for each round

As indicated above, the groups I and II trade in sequence. In odd rounds (1-3-5-...), group I trades good A (for 200 seconds), and then group II trades good B (for 200 seconds). In even rounds (2-4-6-...), group II trades good A (for 200 seconds), and then group I trades good B (for 200 seconds).

The sequence of activities in round 1 will be the following:

1. Group I participants receive information on the value of good A .
2. Group I participants trade good A for 200 seconds, while participants in the other group (II) only observe. While observing the behavior of Group I participants, Group II participants will indicate, on a separate form, what they think the chance is that the true value of good A is 100.
3. Group II participants receive information on the value of good B .
4. Group II participants trade good B for 200 seconds, while participants of the other group only observe.
5. All participants receive information on the outcomes of their trading activity. In particular, everyone will observe the true value of good A and of good B and will be able to compute his/her own payoff according to the rules indicated below.

After the first round is concluded, we start the second round of the experiment. The procedures are identical to those of the first round, with the exception that now group II starts and trades good A and then group I trades good B . The experiment continues until the 10th round is completed.

Trading

The trading platform on your computer, the initial endowment, and the way you sell or buy a good are all the same as in Phase *I*. The only difference is that now the two groups trade different goods (with possibly different values) one after the other, with one group observing the trading activity of the other.

Payment at the end of the experiment

The per-round payoffs will be determined in the same way as in Phase *I*. At the end of every round, you will be told the true value of both goods *A* and *B*. Your total per-round payoff depends on:

- 1) what you hold at the end of the round: the amount of liras plus the value of the units of the good *A* or *B* (according to the good that you traded);
- 2) an extra payment of 5 liras for the first 5 buy or sell trades that you execute (up to a maximum of 25 liras in each round).

We will randomly select 3 out of the 10 rounds of Phase *II* and we will sum your per-round payoffs in these three rounds to determine your final payoff in liras for Phase *II*. We will then convert liras into pounds at the exchange rate of 100 liras = £1 and we will sum up this amount to the participation fee of £5. We will pay you in private, immediately at the end of the experiment.

You will now go through a short questionnaire to make sure that you have understood the instructions and then Phase *II* will start.

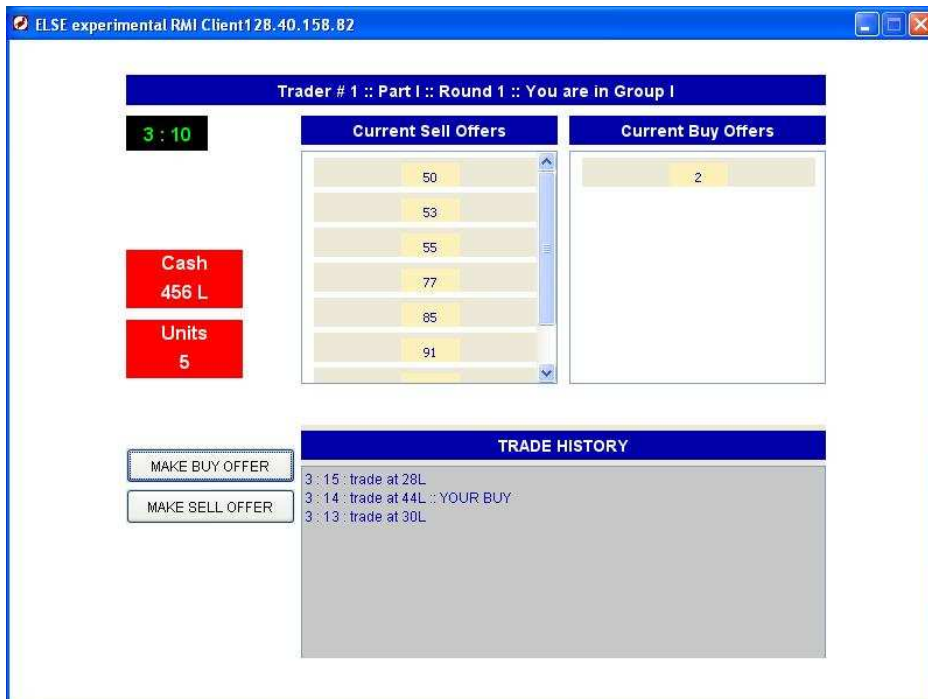


Figure F1: Trading Platform