DEMS
Department of Economics, Management and Statistics

DEMS WORKING PAPER SERIES

When do prediction markets return average beliefs? Experimental evidence

Marco Mantovani, Antonio Filippin

No. 532 – March 2024

Department of Economics, Management and Statistics
University of Milano – Bicocca
Piazza Ateneo Nuovo 1 – 2016 Milan, Italy
http://dems.unimib.it/

When do prediction markets return average beliefs? Experimental evidence ☆

Marco Mantovani[a], Antonio Filippin[b,c]

[a] *University of Milan Bicocca, Department of Economics, Piazza dell'Ateneo Nuovo 1, 20126, Milan*
[b] *University of Milan, Department of Economics, Via Conservatorio 7, 20122 Milano, Italy*
[c] *Institute for the Study of Labor (IZA), Schaumburg-Lippe-Str. 5-9, 53113 Bonn, Germany*

Abstract

In prediction markets prices can be interpreted as the average belief of the traders under restrictive theoretical assumptions, namely specific risk preferences (e.g., log utility) and the prior information equilibrium. Prior information equilibrium is more likely to hold in a call auction, but prediction markets are usually implemented in double auctions that are known to better aggregate information. In this paper we present a laboratory experiment meant to shed some light on this tension, assessing the influence of the main elements that should affect the equilibrium price also manipulating the market institution. We do not find that risk preferences and incorrect beliefs play a significant role in our data. We find instead that in the double auction– where at least partial information aggregation through belief revisions should be expected – prices are closer to the average belief than in the call auction – where, instead, belief aggregation should be expected. We also find evidence that beliefs are updated in the direction of observed prices, rather than of the true state.

JEL Classifications: D81; C92; D40

Keywords: Prediction markets, Information aggregation, Laboratory experiment, Risk preferences, Beliefs

---

☆We would like to thank Peter Bossaerts, Dan Friedman, Anna Gilbert Rivas, Stephanie Heger, Peter Katuščák, Thomas Kittsteiner for useful suggestions. All remaining errors are ours. We gratefully acknowledge financial support from IFREE Small Grant Program and PRIN grant no. P20228SXNF of the Italian Ministry of University and Research.

*Corresponding author: University of Milan, Department of Economics, Via Conservatorio 7, 20122 Milano, Italy
Contact: marco.mantovani@unimib.it (Marco Mantovani), antonio.filippin@unimi.it (Antonio Filippin)

1. Introduction

Conveying information through prices is one of the crucial features of markets. Some exchange systems, the so-called prediction markets, have the aggregation of dispersed beliefs as their primary goal (Arrow et al., 2008; Wolfers and Zitzewitz, 2004). Prediction markets are commonly used as a forecasting device for a wide range of outcomes (see for instance Luckner et al., 2008, and references therein): economic indicators (Gürkaynak et al., 2005), climate outcomes (Hsu, 2011), innovations (Lindič et al., 2011), movie revenues (Gruca et al., 2008), the replicability of results in social sciences (Camerer et al., 2018; Eitan et al., 2018; Forsell et al., 2018), as well as political elections (Berg et al., 2008) and sport events (Spann and Skiera, 2009).

Some support for the accuracy of these forecasts has been found in the field. The Iowa Electronic Market has outperformed opinion polls in the U.S. Presidential elections from 1988 to 2000 (Berg et al., 2008). Using three seasons of the Bundesliga, Spann and Skiera (2009) find that prediction markets and betting odds perform equally well, and both methods strongly outperform tipsters. Laboratory evidence also tends to support the accuracy of prediction markets (see Horn et al., 2014, for a review) at least in simple environments (Healy et al., 2010). This literature also focuses on information acquisition (Page and Siemroth, 2017, 2021), possible incentives to manipulate the outcomes (Choo et al., 2022; Deck and Porter, 2013) and the bias induced by different partitions of the state space (Sonnemann et al., 2013).

Prediction markets aggregate and convey information on beliefs through the price of some state-dependent contracts. The paradigmatic contract is the winner-takes-all (WTA) share, which pays a fixed amount if an event occurs and nothing otherwise. For instance, in the 2016 Brexit referendum the Exit share paid 1 dollar ex post, while Remain paid zero. The market price of this type of asset before the resolution of uncertainty is interpreted as the expected probability that the target event occurs – i.e. the average belief of the traders.

The theoretical conditions under which the price equals the average belief of the traders in a prediction market are well-known in the literature (Gjerstad, 2005; Manski, 2006; Wolfers and Zitzewitz, 2006). At the same time they are rather restrictive. First, the traders must not revise their beliefs inferring information from price data. Second, the traders must be expected utility maximizers with specific risk preferences. The goal of the paper is to study these conditions in a laboratory experiment, and assess their contribution in explaining the differences between market prices and average beliefs. We manipulate *i)* the possibility to make inferences from prices by implementing different market institutions and *ii)* the degree of risk aversion in each market by matching subjects according to their elicited risk preferences. We also assess the role played by failures of Bayesian belief updating in explaining departures of market prices from the average Bayesian belief.

Theoretical models assume that traders in prediction markets hold exogenous beliefs (opinions) about the state of the world on which they agree to disagree (Ottaviani and Sørensen, 2015; Wolfers and Zitzewitz, 2004). For a market to aggregate beliefs as opposed to information, traders must behave according to their ex ante beliefs without extracting further information from (observed or hypothetical) prices. In contexts where beliefs are related to the state of the world by some signal-generating technology, this assumption leads to the Prior Information Equilibrium (PIE), i.e. a Walrasian equilibrium with price-taking traders that do not infer information from prices. Conversely, if traders infer the information possessed by others from (observed or hypothetical) prices, and update their beliefs accordingly, the price no longer represents the traders' beliefs before they enter the market. In the limit case, the process of information extraction would tend to a Rational Expectation Equilibrium (REE). We compare the performance of two market institutions – the single call market and the double auction – as a mechanism for private belief aggregation.

These two institutions induce an exogenous variation in the possibility that traders conform to the PIE assumption. Experimental single call markets usually satisfy the assumption that traders do not anticipate the information conveyed by hypothetical market-clearing prices (Biais et al., 2017; Filippin and Mantovani, 2022; Ngangoué and Weizsäcker, 2021). Recent evidence has shown that the double auction often fails to achieve REE (Choo et al., 2017; Corgnet et al., 2023). At the same time, observing trades and prices helps traders to make inference from prices and revise their beliefs accordingly (Ngangoué and Weizsäcker, 2021). Revised beliefs consider the true state more likely, and their feedback on market activity drives prices in the same direction. Indeed, while both institutions are to some extent inefficient in aggregating information, experimental double auctions outperform call markets.[1] Consistently, we hypothesize that the call auction is the institution that better fits the theoretical conditions for belief aggregation, and that prices will be closer to the true state in the double auction due to (partial) information aggregation. We also

[1]See among others Chen and Plott (2008); Friedman (1984); Friedman and Ostroy (1995); Guarnaschelli et al. (2003); Kagel (2004); Kagel and Levin (1986); Ketcham et al. (1984); Palan et al. (2020); Sakurai and Akiyama (2017); Smith et al. (1982).

gather a direct measure of belief revision based on market activity by eliciting the traders' ex-post beliefs, i.e. after the market closes.

As mentioned above, the risk preferences of the traders also affect belief aggregation in prediction markets (Manski, 2006; Wolfers and Zitzewitz, 2006). Assuming that all traders are characterized by constant relative risk aversion (CRRA) preferences, the equilibrium price reflects the average belief only under logarithmic utility (CRRA coefficient equal to 1). More generally, for low levels of risk aversion less likely events are overpriced relative to the average belief, and the converse hold for high levels (Fountain and Harrison, 2011; Gjerstad, 2005; He and Treich, 2017).

We elicit the participants' risk aversion through an independent task, and use their choices to form markets with substantially different risk-aversion levels (see Crockett et al., 2021, for a similar approach). This exogenous manipulation of the level of risk aversion should have unambiguous effects across markets. Prices should be *a)* further away from the uninformed prior in the more risk-averse markets and *b)* closer to the average beliefs when the CRRA coefficient is closer to 1 (see Section 2 below).

Taken separately, both the underlying theoretical framework and the evidence on (partial) information aggregation in the double auctions are well-established in the literature. Conversely, an empirical test bridging the two things constitutes a novelty of our design and provides an original contribution to the literature. Assessing if and how the market institution affects the accuracy of prediction markets is particularly relevant because real-world prediction markets are implemented as double auctions, while the underlying prediction that prices return the average beliefs ($p^* = \bar{b}$) is derived under assumptions that are less likely satisfied under this institution. Belief revision may drive prices away from the average belief whenever a PIE would instead do so. Moreover, the condition $p^* = \bar{b}$ under PIE also relies upon the assumption of log utility. An empirical exercise comparing the two institutions requires to take risk aversion into account, something that we cannot typically do in the field and, to the best of our knowledge, has not been done before in the lab.

Contrarily to the theoretical predictions, we do not find evidence that risk aversion affect market prices in a significant manner under either market institution.[2] We also find that beliefs before

[2] In call markets, the traders submit their full demand schedules. This allows to study directly if and how risk aversion is expressed in the traders' demands. We carry out this exercise in Filippin and Mantovani (2022) while in this paper we focus on the effect of risk aversion on market prices.

the trading period do not significantly differ, on average, from the Bayesian ones. Moreover, imperfect belief updating do not correlate with deviations of the prices from the average Bayesian beliefs.

The trading institution plays a significant role, with market prices that are closer to the true state in the double auction. While this result per-se is not surprising, the comparison in terms of belief aggregation is counter-intuitive. Prices are closer to the average belief in the double auction than in the call auction. In other words, prices reflect the average amount of information distributed to the market in the institution where the PIE assumption is likely violated, i.e. in the double auction, while they are instead consistently closer to the uninformed prior in the call market, where the PIE assumption is expected to hold.

This difference is not detected at the beginning of the trading period, but only at the end, as double-auction prices converge towards the average beliefs. Surprisingly, such a convergence occurs while ex-post beliefs tends to be revised away from the true state of the world, particularly so for the most informative signals. Beliefs after the trading period show no evidence of sophisticated inferences from prices to the state (see the similar results in Choo et al., 2017). Conversely, beliefs are revised naïvely toward the observed prices. Differences of equilibrium prices across institutions should therefore be attributed to the effect that observing trades and orders has on market activity, rather than on the beliefs themselves. At the same time, the variance of beliefs across traders with different information tend to shrink, i.e. traders disagree less after the market than before. Prices and beliefs co-move in a direction that suggests they eventually converge somewhere in between the uninformed prior and the true state of the world, where also the average belief is located.

The structure of the paper is as follows. Section 2 derives the testable implications. Sections 3 and 4 describe the experimental design and the procedures, respectively. Results are reported in Section 5. Section 6 discusses further implications, while Section 7 concludes.

2. Theoretical background and testable implications

A prediction market is a common-value asset market with heterogeneous beliefs. In this section, we derive our main testable implications by adapting to our experimental setting the standard framework for studying prediction market prices and their relation to average beliefs (Gjerstad, 2005; He and Treich, 2017; Manski, 2006; Wolfers and Zitzewitz, 2006).

There are two ex-ante equally likely states, $e \in \{Red, Blue\}$. An Arrow-Debreu security pays

100 to its owner if $e = Blue$, and pays 0 if $e = Red$. The price of the security is $p \in [0; 100]$. Trader i holds a belief $b_i \in [0, 100]$, representing his subjective probability that $e = Blue$. This percent probability coincides with the expected value of the asset for i. We refer to $b_i = 50$ as to the uninformed belief (or, later on, as to the uninformed prior).

Traders enter the market with an equal endowment m of a numeraire good, which we refer to as 'monetary endowment'. There are no endowments of the security, so that there is no aggregate risk in the market. Sellers take short positions that are covered at the closure of the market at the realized value of the security: 100 if $e = Blue$, 0 if $e = Red$.

We focus on the relation between the market-clearing price p^* and the average belief $\bar{b}$. As shown in Wolfers and Zitzewitz (2006), $p^* = \bar{b}$ when traders' preferences are represented by a constant relative risk aversion (CRRA) utility function with relative risk-aversion coefficient $\theta = 1$. A further assumption is the Prior Information Equilibrium (PIE). That is, traders are price-takers that do not infer others' private information from observed or hypothetical market activity. In this case, the optimal demand of a risk-averse trader is:[3]

$$q_i^*(p, b_i, \theta > 0) = \frac{(1-p)^{\frac{1}{\theta}} b_i^{\frac{1}{\theta}} - p^{\frac{1}{\theta}} (1-b_i)^{\frac{1}{\theta}}}{(1-p)\, p^{\frac{1}{\theta}} (1-b_i)^{\frac{1}{\theta}} + p\,(1-p)^{\frac{1}{\theta}} b_i^{\frac{1}{\theta}}} m, \tag{1}$$

where a negative demand means being a net (short) seller. The optimal demand is equal to zero when $p = b_i$, which is true for any expected utility maximizer, and decreases monotonically as p increases, which is true in general under non-increasing absolute risk aversion. When $\theta = 1$, Equation 1 reduces to:

$$q_i^*(p, b_i, \theta = 1) = (b_i - p)\, \frac{m}{p}. \tag{2}$$

Imposing the market clearing condition $\sum_i q_i^* = 0$ straightforwardly returns a Walrasian equilibrium in which $p^*_{\theta=1} = \bar{b}$.[4]

As pointed out by several contributions, this equality does not hold for other CRRA coefficients or other specifications of the utility function (Gjerstad, 2005; He and Treich, 2017; Manski, 2006). In particular, within CRRA, when risk-aversion is higher than under log-utility ($\theta > 1$), the market price is closer to the true state of the world than the average ex-ante belief. When risk aversion is

[3]See Gjerstad (2005) for a derivation. Risk-neutral and risk-loving players ($\theta \leq 0$) invest their entire endowment in long (short) positions whenever the price is below (above) their belief.

[4]For the sake of simplicity in what follows we do not use an explicit index to identify each market. p^* and $\bar{b}$ should nevertheless be interpreted as market specific variables.

lower ($\theta < 1$), the market price is closer to the uninformed belief, where each state is assigned a 50 percent probability.[5]

This holds also under constant absolute risk aversion (CARA Filippin and Mantovani, 2022) and, more generally, under non-increasing risk aversion if the distribution of beliefs is symmetric (He and Treich, 2017). In these cases, the precise degree of risk-aversion at which $p^* = \bar{b}$ depends on the specific context. Nevertheless, the common pattern is that, as risk aversion increases, prices move away from the uninformed belief in the direction of the true state of the world. This leads to our first testable implication, stating that equilibrium prices are closer to the uninformed belief (50) in less risk-averse markets than in high risk-averse ones:

- **Hp(Risk)** : $\left|p^*_{HIGH} - 50\right| > \left|p^*_{LOW} - 50\right|$,

where p^*_{HIGH} (p^*_{LOW}) is the market-clearing price in a market with High (Low) risk-aversion.

The intuition behind this result is the following. By inducing different demand schedules, risk aversion also affects market prices. The individual net demand (in absolute terms) increases as the degree of risk aversion decreases ($\partial \left|q^*_i\right| / \partial\theta < 0$), ceteris paribus. The extreme case is that of risk-neutral or risk-seeking traders, who invest the entire endowment buying (selling) the asset as long as their belief is above (below) the price. In other words, the higher the degree of risk aversion, the larger needs to be the difference between the price and the belief to demand the same quantity of the asset. The aggregation of less sensitive net demands leads to market prices that are more extreme, because a larger movement of the price from the uninformed belief is needed to clear the market.

Up to now we have considered the PIE – a Walrasian equilibrium given exogenous beliefs and price-taking behavior. When beliefs reflect private information, prices will do so as well. Hence, traders may use (observed or hypothetical) price data to revise their beliefs. If taken to the limit, belief revision leads to the efficient market hypothesis: a REE where endogenous beliefs and prices are consistent. We hypothesize that trading institutions differ in their likelihood of violating the prior information assumption.

In the call auction, traders need to anticipate through introspection the informational content of equilibrium prices. In other words, they need to speculate about the opponents' beliefs (and

[5] This pattern replicates the so-called favorite-longshot bias, an empirical regularity according to which unlikely states are over-priced, and likely states are under-priced (Snowberg and Wolfers, 2010).

the corresponding informational content) that would be consistent with any conceivable market-clearing price, but without observing any price action. The evidence suggests that experimental participants behave non-strategically in call markets (Biais et al., 2017; Ngangoué and Weizsäcker, 2021; Pouget, 2007). Using the same data as those presented in this paper, Filippin and Mantovani (2022) show that behavior in the call auction is empirically indistinguishable from that in a random price mechanism where prices cannot convey any informative content. In this sense, the call auction is an institution in which the prior information assumption holds.

In contrast, the prior information assumption is likely violated in the double auction. Recent evidence has shown that double auction markets are informationally inefficient, in the sense that they typically do not reach the REE (Choo et al., 2017; Corgnet et al., 2023). Nevertheless, at least partial convergence toward this benchmark is often observed, and the double auction involves sequential trades that agents are known to use to extract information (Ngangoué and Weizsäcker, 2021).

Prices that are relatively high (e.g. higher than 50) indicate that the traders' beliefs lean in the direction of $e = Blue$. Traders that realize that these beliefs reflect informative signals on the true state will possibly revise their beliefs in the direction of $e = Blue$, and the feedback effect will further increase prices. The converse holds for relatively low prices. If traders revise their beliefs in this way in the double auction (even without imposing REE), while they act as price takers in the call auction, prices will be closer to the true state in the double auction, and closer to the uninformed belief in the call auction. This constitutes our main testable implication on the effect of the market institution:

- **Hp(Institution):** $|p^*_{DA} - 50| > |p^*_{CA} - 50|$.

where p^*_{DA} (p^*_{CA}) is the market-clearing price in the double (call) auction.

Up to now we have implicitly assumed that beliefs correctly reflect the information received by the traders. That is, traders update the prior in a Bayesian way when they receive private information. As noted in Snowberg and Wolfers (2010) under PIE, incorrect belief updating can be confounded with risk aversion. Incorrect belief updating is also relevant if traders extract information from prices, because it affects the mapping between state, beliefs and market price. Therefore, we design an experiment that allows to control for the actual beliefs of the traders. Our last testable implication is meant to assess whether prices deviate from average Bayesian beliefs in the direction of incorrect belief updating. To do so we aggregate the observations at the market level within each institution and test whether a significant correlation between these variables emerges:

- **Hp(Beliefs):** $\rho(\Delta p, \Delta b) > 0$,

where $\Delta p = p^* - \bar{b}$, is the vector of deviations of the price in each market from the corresponding average Bayesian belief, $\Delta b = \bar{b}_e - \bar{b}$ is the vector of differences at the market level between the average elicited belief and the Bayesian one, and ρ is the linear correlation coefficient.

3. Design

Risk elicitation task. At the beginning of the experiment, we elicit subjects' risk preference using the Investment Game (Gneezy and Potters, 1997). In this task subjects have to decide how to allocate an endowment of 200 Monetary Units (MU) between a safe account and a risky investment that yields 2.5 times the amount invested or zero with equal probability. There exists a closed-form mapping between choices in the investment game and CRRA coefficients. The Investment Game is superior to other tasks in scanning risk preferences around log-utility (Crosetto and Filippin, 2016), which, in theory, is especially relevant in our context. In addition, this risk elicitation method and the theoretical model underlying prediction markets have in common that risk-neutral and risk-seeking agents all behave in the same way.[6]

Matching. We divide each session in two markets of 11 participants, using the median level of elicited risk aversion in the session. This maximizes heterogeneity across the Low and High risk aversion markets, and minimizes heterogeneity within each market. Participants know that there are two separate markets in each session, and that they stay in the same market for all the trading periods, but not that the assignment is made according to their risk preferences.

Asset market. Participants act as traders in 12 trading periods. The structure of every trading period is the same. There are 4 urns that differ in the number of blue marbles they contain out of 100. Urn A contains 47 blue marbles; urn B, 49; urn C, 51; urn D, 53. They are informed that each urn is selected with equal probability, and this uninformed prior is common knowledge.

Traders are endowed with 1000 MU and trade an asset called "Majority Blue". If the urn is C or D, the event "the majority of marbles are Blue" realizes ($e = Blue$) and every asset pays the

[6]The duration of the experiment together with concerns about the salience of the incentives induced us to opt for a single elictation of risk preferences despite the drawback in terms of measurement error (Gillen et al., 2019).

Table 1: Signals

	Signal (s)																
	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58
Urn A	-	-	-	-	-	x	-	-	-	-	-						
Urn B			-	-	-	-	-	x	-	-	-	-	-				
Urn C					-	-	-	-	-	x	-	-	-	-	-		
Urn D							-	-	-	-	-	x	-	-	-	-	-
$p(Blue\|s)$	0	0	0	0	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{2}{3}$	1	1	1	1

Notes: The Table reports the distribution of signals given the selected urn. Each signal is in the form "There are s blue marbles in the urn." Symbols "-" and "x" indicate that the *column* signal is sent under the *row* urn; "x" indicates the correct number of blue marbles in the urn. The last row reports the posterior Bayesian beliefs of a subject receiving the signal in the corresponding column.

owner 100 MU at the end of the trading period. If the urn is A or B, the event "the majority of marbles are Blue" does not realize ($e = Red$) and every asset pays the owner 0 MU.[7]

Before every period, the traders receive a private signal (s) about the composition of the urn, in the following form: "There are s blue marbles in the urn." The signal does not differ by more than 5 units from the true number, i.e. $s \in \{x-5, \ldots, x+5\}$ where x is the true number of blue marbles. All the 11 possible signals are assigned to a different trader. Consequently, each of the 11 traders is assigned one of the 11 signals with the same probability, as illustrated in Table 1. For instance, if the selected urn is A (47 blue), the 11 traders will receive one of the signals ranging from 42 to 52. The procedure that generates and distribute the signals is also common knowledge.

Given the signal, Bayesian updating of the prior about the event generates the beliefs $p(e = Blue|s)$ reported in the last row of Table 1. Some signals ($s \leq 45$ and $s \geq 55$) reveal the state of the world with certainty, leading to $p(e = Blue|s) \in \{0, 1\}$. Other signals ($s = 46, 47, 53, 54$) are partially informative, i.e. $p(e = Blue|s) \in \{\frac{1}{3}, \frac{2}{3}\}$. Finally, some signals ($48 \leq s \leq 52$) are uninformative because the posterior $p(e = Blue|s) = \frac{1}{2}$ coincides with the prior. Aggregating $p(e = Blue|s)$ over the 11 traders in each market returns an average beliefs about the event $e = Blue$, which is equal to 28.8% if the urn is A, 40.9% if the urn is B, 59.1% if the urn is C, and 71.2% if the urn is D. These average beliefs are the benchmark against which we compare the equilibrium

[7]Note that the urn selected is deterministically linked to the value of the asset, there is no draw from the urn.

price.

Belief elicitation. In every trading period, we elicit the subjective belief that each urn has been selected a first time after the traders have received the signal, but before the market opens.[8] We then elicit beliefs a second time after the trading period.

Belief reporting is incentivized using the Binarized Scoring Rule (BSR) (Hossain and Okui, 2013). The BSR compares the sum of squared errors of the reported beliefs (normalized between 0 to 1) with a random number $k \in U[0,1]$. If the sum of squared errors is lower than k – i.e. if the subject's beliefs are sufficiently accurate – he earns a fixed prize (200 MU), otherwise he gets nothing. Instead of paying a different amount according to the accuracy of beliefs like the Quadratic Scoring Rule (QSR), the BSR pays a different probability of receiving the higher of two discrete amounts. Since the participant cannot reduce the variance of the outcomes, his optimal choice must always be that of maximizing the likelihood of getting the high amount, which requires providing the most accurate estimate of the probability distribution regardless of his risk preferences. The BSR is isomorphic to the QRS in terms of expected reward, but it allows to elicit beliefs in an incentive-compatible way regardless of the subject's risk attitudes.[9]

Market institution. In every trading period, the traders enter the market with a monetary endowment of 1000 MU in each trading period. Since they have no endowment of assets, sales are implemented via short selling. This implies that the game is zero-sum and there is no aggregate risk in the market.[10] Net short positions at the end of the trading period are covered at the actual value of the asset. Subjects that are net sellers at the closure of the market buy back the assets they have sold for 0 MU if the urn is A or B, or for 100 MU if the urn is C or D. No-bankruptcy is ensured freezing liquidity for the orders assuming the worst case scenario, i.e. a final value of

[8]Directly eliciting the beliefs about the probability of the event $p(e = Blue|s)$ could have been another possibility. However, it would be prone to mistakes when the subjects combine the probability of the simple events (each urn). This is more likely the case for the signals in which one urn only can be excluded (46,47,53,54).

[9]The QSR induces a truthful revelation of beliefs for a risk-neutral subject but it is not incentive-compatible for a more general class of preferences. For instance, a risk averse player has the incentive to report a smoother distribution of probability, giving up something in terms of expected earnings in exchange for a lower variance of the outcomes. The BSR has its disadvantages, too (Danz et al., 2022). However, we consider it the most suitable procedure in a paper where both risk preferences and beliefs are central.

[10]See, e.g., Bossaerts et al. (2013) and Asparouhova et al. (2017) for a different way to eliminate aggregate risk by design.

the asset equal to zero for net long positions and equal to 100 for the net short positions. This 'single asset - homogenous value' set-up is isomorphic to a two-states/two-assets environment (see Corgnet et al., 2023, for a comparison of these different set-ups).

We run this prediction market in every session under either a call auction or a double auction (between-subjects design):

Call Auction (CA). Traders exchange the asset in a Single Market Call Auction. In every period, they have two minutes to independently place limit orders in a closed book in order to buy and sell the asset. As they submit orders, a visual representation of their net demand schedule updates in real time. The equilibrium price is then computed as the price that equalizes aggregate demand and supply, i.e. where the aggregate net demand is equal to zero, maximizing the volume of trades.[11]

Double Auction (DA). Traders exchange the asset in a Continuous Double Auction. In every period, the market is open for three minutes during which subjects may buy and short sell the asset. Limit orders are posted in an open book. Trades are executed automatically when bid and ask orders match, and traders have full information about the exchanges (price and quantities) that take place in their market in real time.

4. Procedures and Payment scheme

The sessions were run between April 2017 and March 2018 at the EELAB of the University of Milan Bicocca. The experimental software was developed using Z-tree (Fischbacher, 2007).

All sessions follow identical procedures. Upon arrival, participants are randomly assigned to cubicles in the lab. We then elicit risk attitudes and assign subjects either to the High or to the Low risk aversion market. They then receive detailed instruction on the rules and the working of the relevant market institution (CA Vs. DA) and answer a battery of quizzes.[12] The reading of the instructions would move on only once all participants have cleared the quiz. For each quiz, we keep track of the number of mistakes made before providing the correct answer as a proxy for the

[11] In case demand and supply are equal for a range of prices the average of these prices is selected. In case demand and supply do not exactly match, some orders may not be executed (in part). Priority in the execution is given to buy (sell) orders with higher (lower) limit price.

[12] The first quiz regards urns and signals; the second, the belief elicitation procedure; the third, limit orders; the fourth, short selling and monetary consequences of order execution; the last one (two) in CA (DA) the working of the market interface. Complete instructions are attached in Appendix A in the Supplementary Materials.

participant's comprehension of the task.

Urns were assigned to trading periods using a pseudo-random procedure. We generate a random order of 12 periods, where each of the 4 urns is used 3 times. We assign the resulting order of urns to half of the sessions of each institution, and the inverted order to the other half.

Within each of the 12 periods, participants first receive their signal and have 30 seconds to report the probability that each urn has been selected. Then the trading period takes place. The software checks that the no-bankruptcy condition is satisfied independent of the underlying state of the world before accepting an order, and returns an error message in case the condition is violated. In the DA, traders see the order books and the realized exchanges in real time. In the CA traders see only their own individual demand in real time, while they are informed of the equilibrium price and their resulting individual portfolio after the closure of the order submission phase. Under both trading institutions, participants are then asked to report again the probability that each urn has been selected, without receiving any feedback about the true state.

At the end of the experiment the computer selects randomly: (i) the outcome of the Investment Game separately for each participant; (ii) one period for all participants in each session to be used for payments of the trading task; (iii) one of the two measures of beliefs for each participant in one period. To exclude the possibility of hedging within each period, subjects know that the period in (iii) must be different from that in (ii). To compute payments of the belief task according to the BSR, one random number is assigned to every participant. Random numbers and periods in (iii) are different for every participant to avoid social comparison effects. To ensure credibility of our procedures, subjects at the end receive detailed information about the distribution in the session of all the random draws made at the individual level as well as the actual urn selected in every period. Subjects are then notified about their earnings, and they fill in a short questionnaire. Finally they receive their compensation in an anonymous manner.

We collect data from 20 sessions, i.e. 440 experimental subjects, equally split between CA and DA. These numbers correspond to 20 independent observations for each institution – 10 High and 10 Low risk aversion markets – with data on 480 trading periods overall, 120 for each urn. The average duration of the sessions was about two hours and the average payment was 15.9 €.

5. Results

In this section we start by comparing the equilibrium prices with the average beliefs in the two trading institutions. We then investigate the effect of risk aversion on prices. Finally, we study the role of beliefs and of their consistency with signals. All non-parametric tests are based on one independent observation per market, i.e. $N = 20$ for each trading institution.

Prices in Call and Double Auction: Informational inefficiency and average beliefs

Figure 1 summarizes the distributions of closing prices in the two institutions, pooling data of High and Low risk aversion markets. For the DA, we define as 'closing price' in every period the average price of the last ten executed trades. In the CA, all trades are executed at the market-clearing price, and we similarly refer to it as closing price. Prices display informational inefficiency, as they are closer to the average Bayesian belief (PIE - dashed lines in the figure) than to the REE (0 in Urn A and B, 100 in Urn C and D). The absolute difference between the closing price and the average Bayesian belief is smaller than that between the closing price and the true value of the asset (REE) 90 percent of the times in the DA, and 97 percent of the times in the CA.[13]

Comparing the two institutions, closing prices are closer to the average beliefs in the DA than in the CA, where they are instead closer to the uninformed prior (50). Table 2 corroborates this result. A battery of Wilcoxon signed-rank tests does not reject equality between closing prices and the average Bayesian belief in the DA, while it finds a significant difference for urns A, C and D in the CA (Panel A). Again with the exception of urn B, a battery of Mann-Whitney rank-sum tests rejects equality of closing prices between the DA and the CA (Panel B, second row).[14]

Prices in the DA also respond more than in the CA when the traders receive additional information that makes the average beliefs more accurate. Indeed, urns A and D distribute more information than urns B and C, and a Mann-Whitney rank-sum test shows that the average absolute difference in prices between high- and low-information urns ($|p_A - p_B|$ and $|p_C - p_D|$) is larger in the DA than in the CA ($U = 2.705$, p-value = .007).

- **Result(Institution):** Equilibrium prices are significantly closer to the uninformed prior in the CA than in the DA, where they are not significantly different from the average Bayesian belief.

[13]It is common to use the absolute differences between prices and REE/PIE when testing for information aggregation (e.g. Choo et al., 2017; Corgnet et al., 2023; Plott and Sunder, 1982). This literature implicitly assumes risk neutrality when computing the PIE. In the context of our research question, it makes more sense to use the average belief, although it makes little difference in our case. Prices are closer to the risk-neutral PIE than to the REE 82 percent of the times in the DA and 95 percent of the times in the CA.

[14]Testing for the difference across CA and DA of the spread between prices and beliefs returns identical results.

Figure 1: Prices in the double and in the call auction

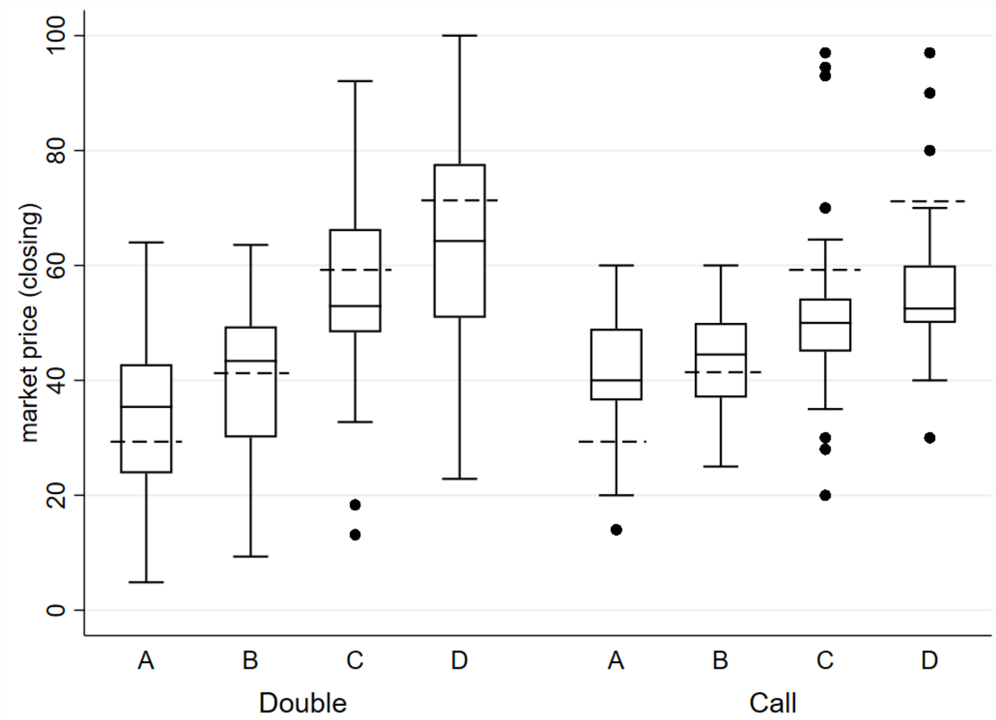

Note: The boxplot shows, for each urn, the distribution of closing prices in the Double Auction and of market-clearing prices in the Call auction. The dashed lines correspond to the average Bayesian beliefs given the urn.

The first part of the main result above is in line with *Hp(Institution)*. The second is not, at least at face value: we conjectured that the CA is the trading institution that better replicates the prior information assumption underlying the interpretation of prices as average beliefs. Or finding may be a composition effect of risk preferences driving prices toward the uninformed prior (in both the DA and the CA) and belief revision (in the DA only). In the following sections we study the role of risk preferences and beliefs. Two things are worth noting before doing this.

First, prices also appear more volatile in the DA according to Figure 1. The larger variance of prices may affect the occurrence of the so-called 'mirages' (Camerer and Weigelt, 1991). A 'mirage' is a situation in which prices are at odds with the information possessed by the market. In our context it corresponds to the instances where the true state of the world is considered as strictly less likely than the wrong one – i.e. $p_A, p_B > 50$ and $p_C, p_D < 50$. Mirages are slightly more frequent (22% against 16%) in the DA. However, the difference is not significant according to a Fisher's exact test (p-value = .297).[15]

[15]On the other hand, in the DA it occurs more often than in the CA that the true state of the world is regarded as

Table 2: Non-parametric tests on prices

Panel A: Within-treatment differences of prices Vs. average Bayesian beliefs by urn								
	Urn A		Urn B		Urn C		Urn D	
	z	p-value	z	p-value	z	p-value	z	p-value
CA	**3.883**	**.000**	1.792	.073	**-3.025**	**.003**	**-3.771**	**.000**
DA	1.717	.086	.149	.881	-1.045	.296	-1.531	.126

Panel B: Across-treatment differences of prices by urn								
	Urn A		Urn B		Urn C		Urn D	
	U	p-value	U	p-value	U	p-value	U	p-value
Opening price	1.271	.204	1.65	.099	.676	.499	.555	.579
Closing price	**-2.083**	**.037**	-.595	.552	**1.956**	**.046**	**2.137**	**.033**

Notes: Panel A reports the Wilcoxon signed-rank test, and corresponding p-value, on the (urn-by-urn) difference between prices and average Bayesian beliefs, separately for the Call ('CA') and the Double ('DA') auction. Panel B reports, for each urn, the Mann Whitney U test statistic, and corresponding p-value, on the (urn-by-urn) difference in prices between the Double and the Call auction. The relevant price is always the market-clearing one in the Call auction. For the Double auction, it is the opening price (first ten trades) in the first row of the panel, and the closing price (last ten trades) in the second row of the panel. A positive statistic means a higher value for the Double auction. Bold indicates significance at the .05 level. All statistics are computed using one observation per market (20 independent observations for each institution).

Second, the difference between the two institutions arises over time. Figure 2 shows that in the DA prices converge towards the average Bayesian beliefs during the trading period. Conversely, 'opening prices' in the DA (mean over the first ten trades) do not significantly differ from the market-clearing prices in the CA, as confirmed by Table 2 (first row of Panel B).

Risk aversion

Figure 3 plots the average degree of risk aversion (CRRA) in each market against the absolute value of the difference between the closing price and the average Bayesian belief in every period, for both the CA and the DA.[16] The CRRA coefficient is estimated for every participant using his choice in the Investment Game. Under CRRA and PIE, this difference should be zero (price equals average beliefs) when the CRRA coefficient is equal to 1, and should increase for

strictly more likely than the wrong one (76% vs 55% of the times, Fisher's exact test, p-value = .001). Prices are more often 'agnostic' ($p = 50$) in the CA than in the DA.

[16]The panels for the CA in Figures 3 and 4 are reproduced from Filippin and Mantovani (2022).

Figure 2: Convergence of prices within the double auction

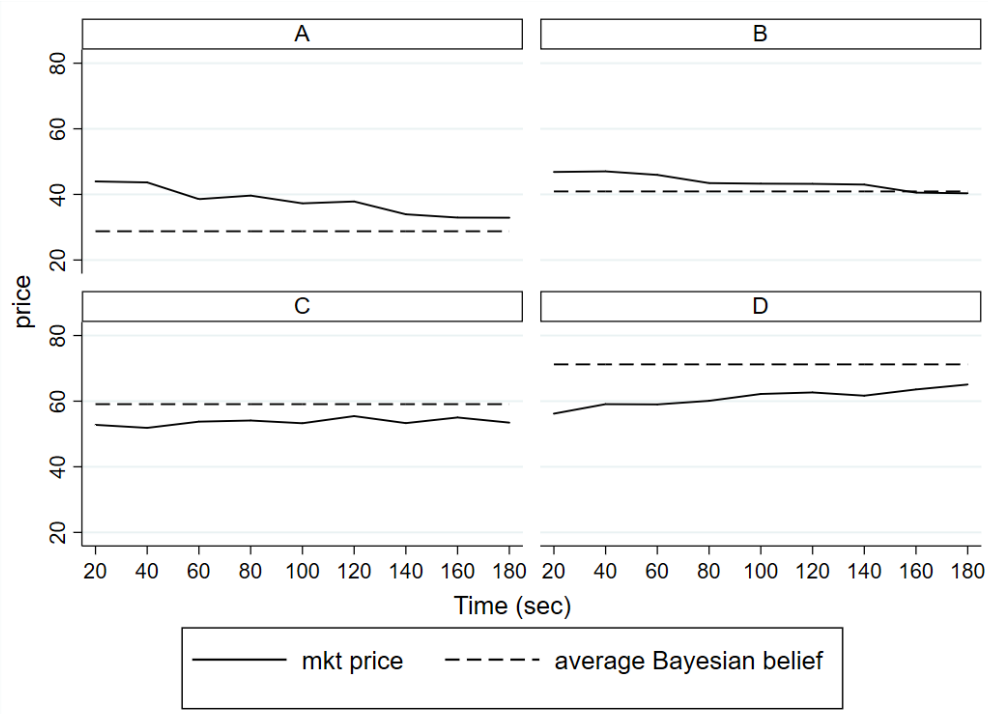

Notes: The Figure shows the evolution of prices – moving average over all markets – over time within a trading period (i.e., 3 minutes) in the Double auction. The dashed lines correspond to the average Bayesian beliefs given the urn.

lower and lower degrees of risk aversion. Therefore, we should expect a negative correlation between the two variables in Figure 3, something that we clearly do not observe. The difference between price and average Bayesian belief is orthogonal to the average degree of risk aversion, and the (non-significant) effect of risk aversion does not significantly differ across institutions. In particular, closing prices are not closer to the average belief when the average CRRA in the market approximate log-utility $(\theta = 1)$.

The matching protocol exogenously manipulates the degree of risk aversion in our markets. We compare the absolute difference between the closing price and the uninformed prior between High and Low risk-aversion markets. A rank-sum test fails to reject the null for both the DA (U = .907, p-value = .364) and the CA (U = −.302, p-value = .762;). Figure 4 shows that the distribution of closing prices is indeed very similar between high and low risk-aversion markets in both the DA and the CA.

- **Result(Risk):** In both institutions prices are not closer to the uninformed prior in Low risk-

Figure 3: Risk aversion and difference between prices and average beliefs

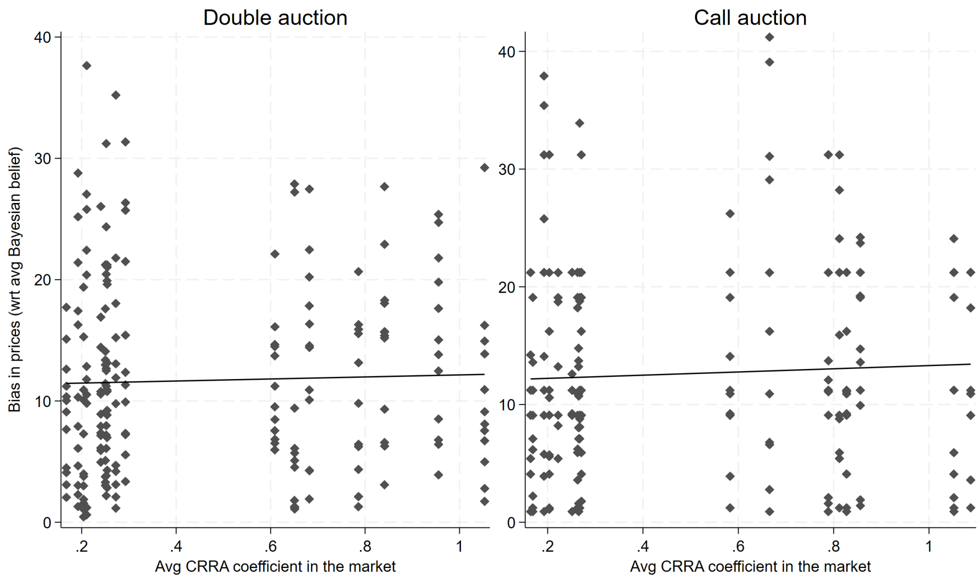

Notes: The Figure displays, for the Double auction, the distance of prices in each market/period from the average Bayesian belief, plotted against the average CRRA coefficient in the market (the gaps between around .3 and .6 are the effect of the matching procedure). A linear fit between the two measures is overimposed.

Figure 4: Risk aversion and prices in the double auction

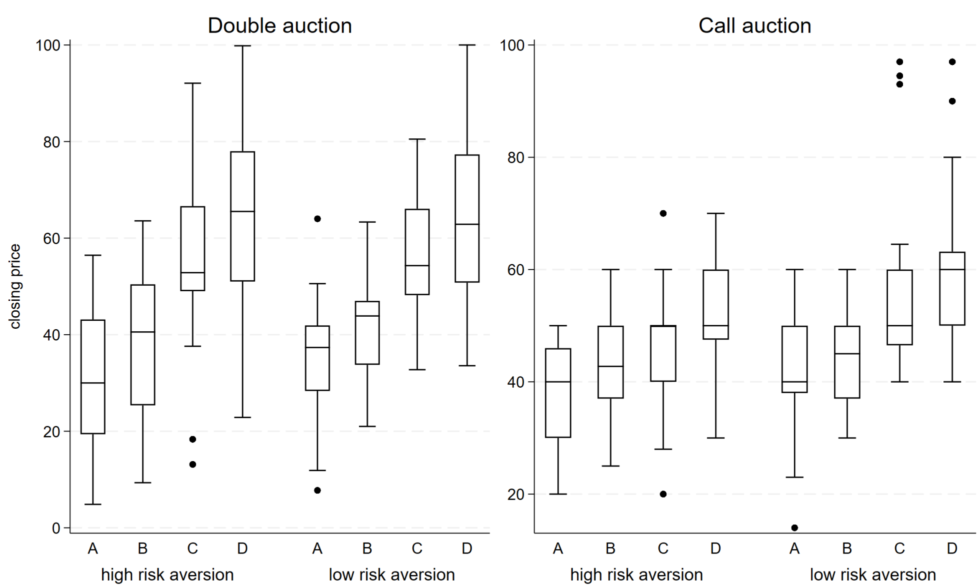

Note: The boxplot shows the distribution of closing prices for each urn under the Double auction, separately for High and Low risk-aversion markets. The dashed lines correspond to the average Bayesian beliefs given the urn.

aversion markets.[17]

[17]Data from the CA, where it is possible to to study directly if and how risk aversion is incorporated in the individual

Table 3: Beliefs

	Signal					Urn			
	42-45	46-47	48-52	53-54	55-58	A	B	C	D
Bayesian	0	33.3	50	66.6	100	28.8	40.9	59.1	71.2
DA ex-ante	4.02	28.25	51.43	76.34	97.60	30.11	42.74	60.11	73.22
CA ex-ante	3.76	27.26	52.54	75.61	97.10	29.46	41.44	61.30	74.29
DA ex-post	13.45	33.02	52.43	69.90	90.33	34.11	44.88	59.84	69.73
CA ex-post	15.99	31.08	49.40	61.35	83.84	33.41	41.19	55.48	64.36

Notes: row 'Bayesian' reports the Bayesian posterior following each type of signal (averaging over urns) and the average Bayesian posterior for each urn (averaging over signals). Each of the rows below reports the corresponding average stated belief at the beginning ('ex-ante') or at the end ('ex-post') of the trading period for each trading instirution (DA: Double Auction, CA: Call Auction).

Bayesian updating and beliefs

Another element that may distort prices away from the average Bayesian belief of the traders is the misperception of the probability that the asset pays off given the signal received, i.e. $p(e = Blue|s)$. Beliefs are obtained by adding the elicited probability assigned to urns C and D. Table 3 reports ex-ante and ex-post beliefs across institutions, signals, and urns.

Ex-ante beliefs trace rather closely the Bayesian ones under both institutions. Not surprisingly, beliefs at the opening of the market do not differ across institutions. Mann-Whitney tests fail to detect a significant difference between CA and DA for any level of the signal.[18] Therefore, individual beliefs cannot explain differences in the prices between the CA and the DA even if they depart from Bayesian ones. Moreover, the distance of ex-ante beliefs from Bayesian ones at the market level does not correlate with the difference between the observed prices and the average Bayesian belief (DA: $\rho = -0.004$, p-value = .948, CA: $\rho = -0.022$, p-value = .739). Therefore, we do not find support for the hypothesis that incorrect Bayesian updating contribute to distort prices.

- **Result(Perception):** Deviations of prices from average beliefs are not explained by incorrect Bayesian updating of beliefs.

demand schedules, are analyzed in more detail in Filippin and Mantovani (2022).

[18]Signals $[42, 45]$: $U = .123$, p-value = .902; $[46, 47]$: $U = -.243$, p-value = .808; $[48, 52]$: $U = -.947$, p-value = .344; $[53, 54]$: $U = .974$, p-value = .330; $[55, 58]$: $U = .950$, p-value = .342.

6. Discussion

In the DA prices at the beginning of the trading period are indistinguishable from those in the CA, but they converge towards the average Bayesian beliefs during the trading activity. Our hypothesis on the effect of the market institution is based on the idea that in the DA traders revise their beliefs after observing prices in the direction of the true state, leading to (at least partial) information aggregation. We can scrutinize this idea because we observe beliefs also at the end of each trading period.

Table 3 shows that ex-post beliefs are indeed revised following the market activity. However, subjects fail to infer the true state of the world. Conversely, ex-post beliefs tends to be revised away from the true state of the world, particularly so for the most informative signals. The fully informed traders (those receiving the signals $42-45$, $55-58$) become more uncertain about the state of the world. The process of beliefs updating is similar in the two institutions, although ex-post beliefs departs from the Bayesian ones to a slightly lower extent in the DA.

More formally, we measure the precision of individual beliefs using the sum of the squared error relative to the true state of the world. For each type of signal, we then test for the difference between the ex-ante and the ex-post sum of squared errors using Wilcoxon signed-rank test. There is no sign that beliefs become on average more precise after trading. In the CA ex post beliefs significantly worsen for almost each type of signal.[19] Even in the DA, beliefs tend to worsen on average, and significantly so for the most informative signals.[20] If we run the same test without conditioning on the signal, a randomly picked individual is found to be more informed about the state of the world before observing any market activity than afterward, under both institutions.[21]

Rather than through sophisticated inferences from prices to the state, the updating process seems to go through naïve revisions in the direction of the observed price. Figure 5 shows, for each urn, signal and institution, the direction of the revision from the average ex-ante to the average ex-post belief. The figure also plots the average closing price and specifies the actual value of the asset representing the true state of the world. Beliefs are revised in the direction of the price virtually in all cases, while the true value of the asset plays little role. Prices are relatively more

[19] Signals $[42,45]$: $=-3.920$, p-value $<.001$; $[46,47]$: $z=-1.792$, p-value $=.073$; $[48,52]$: $z=-2.352$, p-value $=.019$; $[53,54]$: $z=-2.725$, p-value $=.006$; $[55,58]$: $z=-3.920$, p-value $<.001$.

[20] Signals $[42,45]$: $=-3.285$, p-value $=.001$; $[46,47]$: $z=-.336$, p-value $=.737$; $[48,52]$: $z=-.037$, p-value $=.970$; $[53,54]$: $z=-1.680$, p-value $=.093$; $[55,58]$: $z=-3.568$, p-value $<.001$.

[21] DA: $z=-2.800$, p-value $=.005$; CA: $z=-3.733$, p-value $<.001$.

Figure 5: Revision of beliefs after observing market prices

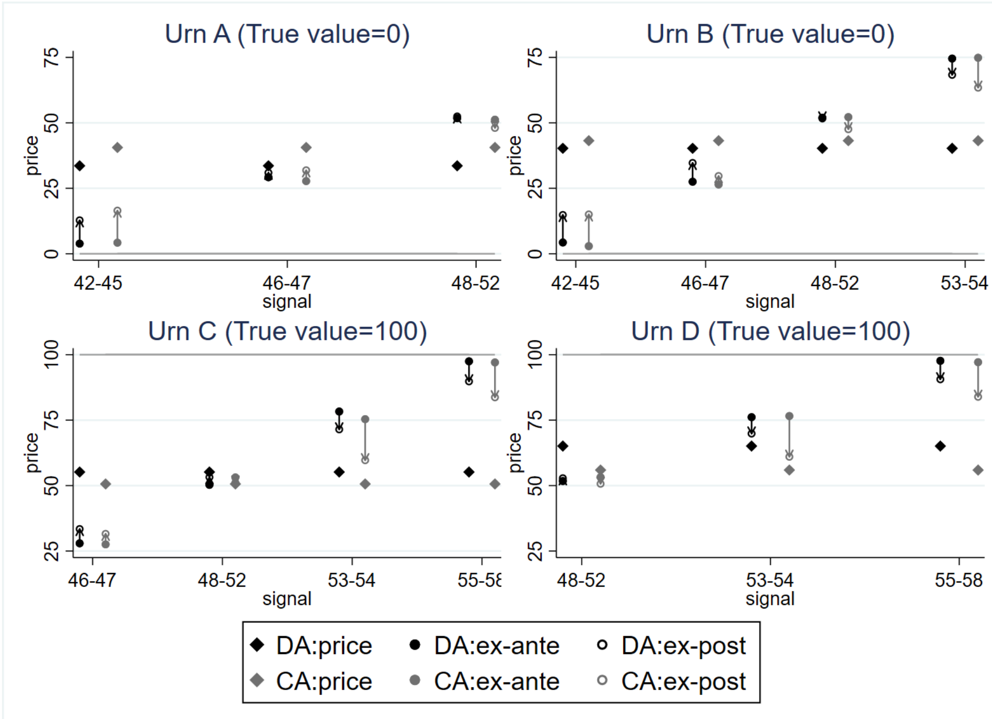

Note: each plot is specific to a urn and provides information on ex-ante (cirlces) and ex-post (hollow circles) average beliefs, in both the DA (black symbols) and the CA (gray symbols). Beliefs are presented for each type of signal distributed under the corresponding urn. Arrows are used to connect ex-ante to ex-post beliefs. Diamond symbols represent the average price under the corresponding urn and institution. The gray line in each plot is the actual value of the asset.

informative about the state in the DA than in the CA. Consistently, ex-post beliefs turn out to also be comparatively more accurate in the DA. In fact, the sum of squared errors of ex-post beliefs is significantly lower in the DA than in the CA (Mann-Whitney rank-sum test: $U = 2.894$, p-value $= .003$).

The knowledge about the state of the world does not improve in the DA where the trading activity is observed. This finding is in line with the literature showing that beliefs do not converge to the REE (Choo et al., 2017). At the same time, it seems to exclude information aggregation as the explanation of the differences between the DA and the CA. Moreover, observing that prices improve while beliefs worsen is definitely counter-intuitive and suggests that those differences must be traced back to the effect that observing trades and orders has on market activity, i.e. effects on how traders behave for given beliefs, rather than on the beliefs themselves. The question is what mechanism might generate these effects.

In the CA, the traders fail to transfer the information they possess into their trading activity.

In fact, their demands pivot (i.e. switch from positive to negative) around a price that is not equal to their belief, as it should under any expected utility specification, but instead closer to the uninformed prior. Filippin and Mantovani (2022) analyze individual demands in the CA and label *operational conservatism* this tendency to act *as if* traders had less information than they do. In the DA, the traders' activity seems to incorporate their private information to a larger extent. However, it is not obvious how to formally retrieve the net demands in the DA in order to compare them to those in the CA.

This explanation may clarify why prices move toward the true state even while beliefs move in the opposite direction and toward the price. Since traders are acting as if they had a belief closer to 50 than their Bayesian belief, their demand can become more consistent with their Bayesian belief even if their belief worsens during the trading period.[22]

There are several possible reasons why traders end up transferring more information into their trading activity in the DA. First, the inability to transfer the information possessed in the trading activity could be related to perceived ambiguity in complex environments (Asparouhova et al., 2015; Dimmock et al., 2016; Fattinger, 2018; Trautmann and Van De Kuilen, 2015).[23] It could then be that traders end up perceiving the environment as less complex/ambiguous during the trading period in the DA. Second, strategic interaction may lead traders to bid more aggressively. As an illustrative example, traders who are sure that the asset will pay may start bidding at low prices, similar to those in the CA, aiming at high per-unit profits. However, observing the orders of their opponents may foster more and more competitive bids, resulting in demands that better incorporate their information and inducing higher market prices.

While our experiment is not equipped to screen these possible explanations, a potential implication of the dynamics of prices and beliefs is worth being stressed. As beliefs get closer to the price, their variance across traders with different information tend to shrink. In fact, the variance across different signals between ex-ante and ex-post beliefs decreases by 25 and 21 percent in urns

[22]For instance, a trader with a $b_i = .04$ (average ex-ante belief in the DA for $s_i \in [42 - 45]$) at the beginning of the trading period may bid consistently with, say, $b_i = .25$. Even if she revise her belief so that $b_i = .13$ (average ex-post belief in the DA for the same signals), her bids at the end of the trading period can still be more in line with the information possessed (that should imply $b_i = 0$) than they were at the beginning.

[23]Similar effects in individual decision problems result in 'complexity aversion' (Bernheim and Sprenger, 2020) or in perceiving compound lotteries as ambiguous (Halevy, 2007). See also Huber et al. (2019) on the effect of different risk perceptions of equally risky assets.

A, D and B, C, respectively.[24] In other words, traders disagree less after the market than before. This fact seems to exclude that prices can ever return the true state of the world. In contrast, the co-movement of prices and beliefs suggests they may eventually converge somewhere in between the uninformed prior and the true state of the world. By construction, the average belief is located there, although there is no guarantee it represents the exact point of convergence.

7. Conclusion

Prediction markets promise a mechanism for aggregating dispersed beliefs that provides proper incentives while being virtually costless. Prediction markets have grown and diffused in several fields despite the absence of a solid empirical validation of the underlying theory. This theory builds on the assumption of price-taking behavior (Prior Information Equilibrium - PIE). It also implies that risk preferences affect the equilibrium price and therefore the interpretation of prices as average beliefs.

We design an experiment where markets differ in the degree of risk aversion of the traders. We manipulate the market institution (call vs. double auction) in order to compare markets where the PIE benchmark is more or less likely to hold. We also elicit the traders beliefs before the market, to control for incorrect Bayesian updating, and after the market, to assess the extent to which subjects infer information from the observed prices.

Our results show that prices are relatively close to the average beliefs in the double auction, while they are close to an uninformed belief in the call auction. While a relatively higher informational efficiency in the double auction is expected, this result is surprising given that the institution that best conforms to the PIE is the call auction. Our design allows to screen risk aversion as an explanation for prices being more informative. However, we do not find evidence that the traders' risk preferences substantially affect market prices.

The crucial element that distinguishes the double auction are feedback effects between observed prices, beliefs, and one's trading activity. While this mechanism seems to facilitate inference about the true state of the world, our data suggest that this is not the case. In fact, beliefs are revised in the direction of the observed prices, rather than of the true state of the world, after the trading period. Nevertheless, individual orders end up transferring more of the information of

[24] The variance of Bayesian beliefs is the same in urns A and D, and so it is in urns B and C, so we analyze each pair together.

the trader into market prices. This evidence suggests that the feedback effects that matter in the double auction go directly from observed market activity to one's demand, rather than through sophisticated inference about the underlying state.

The answer to the question in our title – when do prediction markets return average beliefs? – is not clear-cut. On the one hand, the convergence of prices and beliefs in the double auction looks as promising for prediction markets to serve as a costless machanism of belief aggregation. On the other, call markets fail at the task. At the same time, the average trader ends up being less informed after the market than he was before in both institutions. Directly eliciting beliefs is more costly, but may outperform prediction markets without incurring in undesired spill-overs on the traders' beliefs.

References

Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., Levmore, S., Litan, R., Milgrom, P., Nelson, F. D., et al., 2008. The promise of prediction markets. Science 320 (5878), 877–878.

Asparouhova, E., Bossaerts, P., Eguia, J., Zame, W., 2015. Asset pricing and asymmetric reasoning. Journal of Political Economy 123 (1), 66–122.

Asparouhova, E., Bossaerts, P., et al., 2017. Experiments on percolation of information in dark markets. Economic Journal 127 (605), 518–544.

Berg, J., Forsythe, R., Nelson, F., Rietz, T., 2008. Results from a dozen years of election futures markets research. Handbook of experimental economics results 1, 742–751.

Bernheim, B. D., Sprenger, C., 2020. On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting. Econometrica 88 (4), 1363–1409.

Biais, B., Mariotti, T., Moinas, S., Pouget, S., 2017. Asset pricing and risk sharing in a complete market: An experimental investigation. Tech. rep., Toulouse School of Economics (TSE).

Bossaerts, P., Frydman, C., Ledyard, J., 2013. The speed of information revelation and eventual price quality in markets with insiders: comparing two theories. Review of Finance 18 (1), 1–22.

Camerer, C., Weigelt, K., 1991. Information mirages in experimental asset markets. Journal of Business, 463–493.

Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B. A., Pfeiffer, T., et al., 2018. Evaluating the replicability of social science experiments in nature and science between 2010 and 2015. Nature Human Behaviour 2 (9), 637.

Chen, K.-Y., Plott, C. R., 2008. Markets and information aggregation mechanisms. Handbook of Experimental Economics Results 1, 344–352.

Choo, L., Kaplan, T. R., Zultan, R., 2017. Information aggregation in arrow–debreu markets: an experiment. Experimental Economics, 1–28.

Choo, L., Kaplan, T. R., Zultan, R., 2022. Manipulation and (mis) trust in prediction markets. Management Science.

Corgnet, B., Deck, C., DeSantis, M., Hampton, K., Kimbrough, E. O., 2023. When do security markets aggregate dispersed information? Management Science.

Crockett, S., Friedman, D., Oprea, R., 2021. Naturally occurring preferences and general equilibrium: a laboratory study. International Economic Review 62 (2), 831–859.

Crosetto, P., Filippin, A., 2016. A theoretical and experimental appraisal of four risk elicitation methods. Experimental Economics, 613–641.

Danz, D., Vesterlund, L., Wilson, A. J., 2022. Belief elicitation and behavioral incentive compatibility. American Economic Review.

Deck, C., Porter, D., 2013. Prediction markets in the laboratory. Journal of Economic Surveys 27 (3), 589–603.

Dimmock, S. G., Kouwenberg, R., Wakker, P. P., 2016. Ambiguity attitudes in a large representative sample. Management Science 62 (5), 1363–1380.

Eitan, O., Viganola, D., Inbar, Y., Dreber, A., Johannesson, M., Pfeiffer, T., Thau, S., Uhlmann, E. L., 2018. Is research in social psychology politically biased? systematic empirical tests and a forecasting survey to address the controversy. Journal of Experimental Social Psychology 79, 188–199.

Fattinger, F., 2018. Trading complex risks. Available at SSRN 3086358.

Filippin, A., Mantovani, M., 2022. Risk aversion and information aggregation in binary asset markets. Quantitative Economics, forthcoming.

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

Forsell, E., Viganola, D., Pfeiffer, T., Almenberg, J., Wilson, B., Chen, Y., Nosek, B. A., Johannesson, M., Dreber, A., 2018. Predicting replication outcomes in the many labs 2 study. Journal of Economic Psychology.

Fountain, J., Harrison, G., 2011. What do prediction markets predict? Applied Economics Letters 18 (3), 267–272.

Friedman, D., 1984. On the efficiency of experimental double auction markets. The American Economic Review 74 (1), 60–72.

Friedman, D., Ostroy, J., 1995. Competitivity in auction markets: An experimental and theoretical investigation. The Economic Journal 105 (428), 22–53.

Gillen, B., Snowberg, E., Yariv, L., 2019. Experimenting with measurement error: Techniques with applications to the caltech cohort study. Journal of Political Economy 127 (4), 1826–1863.

Gjerstad, S., Nov. 2005. Risk Aversion, Beliefs, and Prediction Market Equilibrium. Microeconomics 0411002, EconWPA.

Gneezy, U., Potters, J., 1997. An Experiment on Risk Taking and Evaluation Periods. The Quarterly Journal of Economics 112 (2), 631–45.

Gruca, T. S., Berg, J. E., Cipriano, M., 2008. Incentive and accuracy issues in movie prediction markets. Journal of Prediction Markets 2 (1), 29–43.

Guarnaschelli, S., Kwasnica, A. M., Plott, C. R., 2003. Information aggregation in double auctions: Rational expectations and the winner's curse. Information Systems Frontiers 5 (1), 63–77.

Gürkaynak, R. S., Wolfers, J., Carroll, C. D., Szeidl, A., 2005. Macroeconomic derivatives: An initial analysis of market-based macro forecasts, uncertainty, and risk. In: NBER International Seminar on Macroeconomics. Vol. 2005. The University of Chicago Press Chicago, IL, pp. 11–64.

Halevy, Y., 2007. Ellsberg Revisited: An Experimental Study. Econometrica 75 (2), 503–536.

He, X.-Z., Treich, N., 2017. Prediction market prices under risk aversion and heterogeneous beliefs. Journal of Mathematical Economics 70, 105–114.

Healy, P. J., Linardi, S., Lowery, J. R., Ledyard, J. O., 2010. Prediction markets: alternative mechanisms for complex environments with few traders. Management Science 56 (11), 1977–1996.

Horn, C. F., Ivens, B. S., Ohneberg, M., Brem, A., 2014. Prediction markets–a literature review 2014. The Journal of Prediction Markets 8 (2), 89–126.

Hossain, T., Okui, R., 2013. The binarized scoring rule. The Review of Economic Studies 80 (3), 984.

Hsu, S.-L., 2011. A prediction market for climate outcomes. U. Colo. L. Rev. 83, 179.

Huber, J., Palan, S., Zeisberger, S., 2019. Does investor risk perception drive asset prices in markets? experimental evidence. Journal of banking & Finance 108, 105635.

Kagel, J. H., 2004. Double auction markets with stochastic supply and demand schedules: Call markets and continuous auction trading mechanisms. In: Advances in Understanding Strategic Behaviour. Springer, pp. 181–208.

Kagel, J. H., Levin, D., 1986. The winner's curse and public information in common value auctions. The American economic review, 894–920.

Ketcham, J., Smith, V. L., Williams, A. W., 1984. A comparison of posted-offer and double-auction pricing institutions. The Review of Economic Studies 51 (4), 595–614.

Lindič, J., Baloh, P., Ribière, V. M., Desouza, K. C., 2011. Deploying information technologies for organizational innovation: Lessons from case studies. International Journal of Information Management 31 (2), 183–188.

Luckner, S., Schröder, J., Slamka, C., 2008. On the forecast accuracy of sports prediction markets. In: Negotiation, auctions, and market engineering. Springer, pp. 227–234.

Manski, C. F., 2006. Interpreting the predictions of prediction markets. Economics Letters 91 (3), 425–429.

Ngangoué, M. K., Weizsäcker, G., 2021. Learning from unrealized versus realized prices. American Economic Journal: Microeconomics 13 (2), 174–201.

Ottaviani, M., Sørensen, P. N., 2015. Price reaction to information with heterogeneous beliefs and wealth effects: Underreaction, momentum, and reversal. American Economic Review 105 (1), 1–34.

Page, L., Siemroth, C., 2017. An experimental analysis of information acquisition in prediction markets. Games and Economic Behavior 101, 354–378.

Page, L., Siemroth, C., 2021. How much information is incorporated into financial asset prices? experimental evidence. The Review of Financial Studies 34 (9), 4412–4449.

Palan, S., Huber, J., Senninger, L., 2020. Aggregation mechanisms for crowd predictions. Experimental economics 23 (3), 788–814.

Plott, C. R., Sunder, S., 1982. Efficiency of experimental security markets with insider information: An application of rational-expectations models. Journal of political economy 90 (4), 663–698.

Pouget, S., 2007. Adaptive traders and the design of financial markets. The Journal of Finance 62 (6), 2835–2863.

Sakurai, H., Akiyama, E., 2017. Bubbles and information in continuous double auction and call market: An experiment. Available at SSRN 2930083.

Smith, V. L., Williams, A. W., Bratton, W. K., Vannoni, M. G., 1982. Competitive market institutions: Double auctions vs. sealed bid-offer auctions. The American Economic Review 72 (1), 58–77.

Snowberg, E., Wolfers, J., 2010. Explaining the favorite–long shot bias: Is it risk-love or misperceptions? Journal of Political Economy 118 (4), 723–746.

Sonnemann, U., Camerer, C. F., Fox, C. R., Langer, T., 2013. How psychological framing affects economic market prices in the lab and field. Proceedings of the National academy of Sciences 110 (29), 11779–11784.

Spann, M., Skiera, B., 2009. Sports forecasting: a comparison of the forecast accuracy of prediction markets, betting odds and tipsters. Journal of Forecasting 28 (1), 55–72.

Trautmann, S. T., Van De Kuilen, G., 2015. Ambiguity attitudes. The Wiley Blackwell handbook of judgment and decision making 1, 89–116.

Wolfers, J., Zitzewitz, E., 2004. Prediction markets. Journal of economic perspectives 18 (2), 107–126.

Wolfers, J., Zitzewitz, E., 2006. Interpreting prediction market prices as probabilities. Tech. rep., National Bureau of Economic Research.

Appendix A. Experimental instructions

Welcome to this experiment and thank you very much for taking the time to support our research. In the next two hours you will perform several tasks that are explained in due course. It is a standard practice in this type of studies to provide written instructions to participants and to read them aloud, to ensure that everyone receives the same information.
During the whole experiment, the amounts are expressed in Monetary Units, called MU, whose unit value is one euro cent, so 100 MU = 1 euro.

Investment Game

You have an endowment of 200 MU and you have to choose the portion of this amount (from 0 to 200) that you want to invest in a risky option. Non-invested money directly enters your final earnings.
There is a 50% chance that the investment in the risky option will be successful. If successful, you receive 2.5 times the amount invested. If the investment fails, you lose the amount invested.
The outcome of the risky option will be determined at the end of the experiment flipping a virtual coin:

- If Head shows up the investment is successful and you receive the amount not invested plus 2.5 times the investment;
- If Tail shows up the investment has not been successful and you receive only the amount not invested.

The computer will determine the outcome separately for each of you. To ensure the fairness of the coin toss, everyone will be shown the distribution of outcomes (Head and Tail) in the whole session.

[PLAY THE INVESTMENT GAME]

Market

Your task in the market is to exchange an asset on a computer based trading system. Your compensation in this phase depends on your performance on the market, so listen the instructions carefully. There will also be questions to verify your understanding and you need to provide the correct answer to proceed. If you are told that your answers are wrong and from the error message you do not understand why, please raise your hand. One of us will come to your cubicle to dispel your doubts privately.

Each market is composed of 11 traders (i.e. there are 2 markets in this session). The assignment to one of the two markets takes place at the beginning and lasts for the whole experiment. The experiment consists of 12 trading periods of two 3 minutes **[CA:** 2 minutes**]** each.

Let's now answer in detail to the following questions:

1. What is the setting?
2. What is traded?
3. How does the trading system work?
4. How are your earnings determined?

1. What is the setting?

There are four urns containing red and blue marbles in the following proportions:

- Urn A: 47 red marbles 53 blue marbles
- Urn B: 49 red marbles 51 blue marbles
- Urn C: 51 red marbles 49 blue marbles
- Urn D: 53 red marbles 47 blue marbles

You will know which is the urn actually used in each period only at the end of the experiment. You know that each urn has the same probability of being selected in each period.
Before the market opens in every period, each of you will receive an inaccurate signal about the composition of the urn actually selected. The signals for the different urns are:
You receive with the same probability one of the possible signals given the selected urn. Signals are randomly assigned and are different for each participant in the same market.
As you can see, it is very unlikely that the signal received exactly matches the number that identifies the urn. Nevertheless, the signal can help you to understand if some urn has been selected or not, and this as we shall see is very important to know how to operate on the market in order to maximize your profits.

	Signal																
	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58
Urn A	-	-	-	-	-	x	-	-	-	-	-						
Urn B			-	-	-	-	-	x	-	-	-	-	-				
Urn C					-	-	-	-	-	x	-	-	-	-	-		
Urn D							-	-	-	-	-	x	-	-	-	-	-

[QUIZ 1]

2. What is traded?

This market exchanges an asset called "Majority Blue." When the urn selected is either A or B the event "highest number of blue marbles" does not occur and any asset hold at the end of the trading period is worthless (0 MU). Conversely, if the urn selected is C or D, the event occurs and any asset pays the owner 100 MU.

The value of the asset therefore depends on the urn selected. You have the signal about the urn to think about which value you attribute to the asset, so as to decide how much you are willing to pay to buy the asset and at how much you are willing to sell it on the market. Since the final value of the asset is uncertain, the price at which assets are traded may vary between the two extremes (0 and 100).

Your expectations on the urn selected are therefore essential to guide your choices and determine your earnings. For this reason, in each of the 12 trading periods you will be asked twice what is in your opinion the probability that each urn has been drawn: the first time before the trading period, after receiving the signal, the second time at the end of the trading period.

Expectations

You will be asked to assign an integer number between 0 and 100 to each urn. Such a number represents your estimate of the probability that each urn has been selected. 0 means "certainly not selected" and 100 "certainly selected". The sum of the percentages must be 100.

You can receive an additional compensation based on the accuracy of your expectations. An error index going from zero (perfect estimate) to 100 (completely wrong estimate) will be calculated. The exact formula of the index is complex, and we are happy to explain it at the end of the experiment to those interested in. At the moment it is enough to know that the error index:

- is equal to 0 assigning all the probability to the urn actually selected;

- is equal to 37.5 assigning the same probability to all the urns;

- is equal to 100 when assigning all the probability to a wrong urn.

At the end of the experiment one of the estimates will be chosen randomly. We will then extract a number between 0 and 100 (and to guarantee that periods and numbers are chosen randomly we will show you the distribution drawn in the whole session). If your error index is lower than the selected number it means that your estimates are sufficiently accurate and you will receive 200 MU. If the index is higher, you will not receive any additional compensation.

You have 30 seconds to enter your expectations, after which the experiment proceeds automatically. If you did not enter your expectations in the round relevant towards your earnings you will not receive the 200 MU. To maximize the probability of receiving the 200 MU you must minimize the error index, making the best possible estimate. Given these incentives it is impossible to increase the probability of receiving the 200 MU by distorting the estimate of the probabilities that you have in mind.

Practical advice on how to assign the probabilities:

- if you believe that an urn has been selected with higher likelihood assign it a higher percentage;

- if you believe that two or more urns have been selected with the same likelihood assign them equal percentages;

- allocate all the probability to one or two urns only if you are sure that the selected one is among them.

- do not concentrate the probability on one or two urns if you are not confident that the selected one is among them. If the selected urn is another one you will not earn the 200 MU.

- always report your expectations: a wrong estimate is in any case better than nothing.

[QUIZ 2]

3. How does the trading system work?

[DA: The trading system is a so-called continuous double auction, which we will explain in detail later, whose role is to manage buy and sell orders. Two orders will be executed against each other if the buy order has a limit price greater than or equal to the limit price of the sell order.

[QUIZ 3, DA]

At the beginning of each period you receive an endowment of 1000 MU. You cannot transfer MUs from one period to another. During the market activity, you can enter both buy and sell orders.

Buying

To buy the desired amount of assets, it is sufficient to use the available liquidity and to cross a suitable sell order. The buyer makes profits if he sells the assets at a higher price, or if the final value of the assets is higher. Conversely, a purchase involves a loss if the selling price or the final value of the asset is lower than what has been paid to buy.

Selling

If you buy assets in the market, you can sell them at any time. However, at the beginning of the period you have 1000 MU in your account, while you do not have an endowment of assets. How can you sell assets in this case? It is possible, through the so-called short selling.
Short selling consists in selling assets that you do not possess, as if you borrow them, committing to their subsequent repurchase (a.k.a. covering the short position).
Covering can take place in two ways: a) in the market, during the exchange period, at the price at which the assets are subsequently repurchased; b) at the final value of the asset (0 for urns A and B, 100 for urns C and D), in case assets are still held "short" at the end of the exchange period.
When a sell order is executed, the seller makes profits if the cost of the repurchase is lower than the amount received with the initial short selling. On the other hand, short selling involves a loss if the final value of the asset is higher at the time of repurchase than the price received for its sale.**]**
[CA: In every period the market stays open for 2 minutes after which the trading system computes the market price by combining buy and sell orders. The market price is the price that maximizes the volume of exchanges by matching the quantity purchased and the quantity sold (more on this later).
Once the market price has been computed, the system executes at that price:

1) buy orders issued with a limit price greater than or equal to the market price
2) sell orders issued with a limit price that is less than or equal to the market price.

N.B. The limit price determines whether an order is executed or not, but does not represent the price at which the exchange takes place. All exchanges take place at the market price.

[QUIZ 3, CA]

At the beginning of each period you receive an endowment of 1000 MU. You cannot transfer MUs from one period to another. During the market activity, you can enter both buy and sell orders.

Profits of the buyer

When a buy order is executed, the buyer makes profits if the final value of the asset is higher than the market price paid to buy it. Conversely, a purchase results in a loss if the final value of the asset is lower than the price paid to buy it.

Profits of the seller

At the beginning of the period you have 1000 MU in your account, while you do not have an endowment of assets. How can you sell assets in this case? It is possible, through the so-called short selling.
Short selling consists in selling assets that you do not possess, as if you borrow them, committing to their subsequent repurchase (a.k.a. covering the short position). The repurchase takes place at the final value of the asset (0 for urns A and B, 100 for urns C and D).
When a sell order is executed, the seller makes profits if the cost of the repurchase is lower than the amount received with the initial short selling. On the other hand, short selling involves a loss if the final value of the asset is higher at the time of repurchase than the price received for its sale.**]**

Note that in this market, buying and short selling are symmetric. Since the price of the asset is limited between 0 and 100 it is not possible to make unlimited losses, contrarily to what may happen in the stock exchange.
Consider the following example: in a market a single exchange of 10 assets at a price of 50MU occurs. At the end of the trading period the two traders involved will have the following situation:**]**

	Liquidity	**Assets**	**Final value of the assets**		**Total earnings**	
			Urn A or B	**Urn C or D**	**Urn A or B**	**Urn C or D**
Buyer	500	10	0	1000	500	1500
Seller	1500	-10	0	-1000	1500	500

If the selected urn is A or B the value of the asset is zero: the seller can cover his short position for free, and will have 1500 MU in his account. The buyer has 10 worthless assets, and 500 MU in his account.
If the selected urn is C or D the value of the asset is 100MU. In this case the seller must spend 1000MU to cover the short position, so he will have 500MU left in his account. The buyer holds 10 assets worth 100MU each, so he will have 1500 MU in his account at the end.
As you can see, the two situations are symmetric (remember that each urn has the same probability of being drawn).
Practical advice: If you think the urn selected is A or B (final value of the asset = 0) you should sell short and hold a negative number of assets. If you think the selected urn is C or D (final value = 100) you should buy and hold a positive number of assets instead. Therefore, insert buy and sell orders with the same simplicity: buy when you think the asset is worth 100, sell short when you think the asset is worth zero. Remember that you make profits selling at a higher price than you paid to buy. Therefore, it makes no sense to enter buy orders with a limit price higher than that of a sell order of yours.

[QUIZ 4]

Liquidity

All orders must have financial coverage and for this reason some liquidity is frozen when orders are submitted. Freezing liquidity ensures that at any market price the execution of all the pending orders does not require more than the 1000 MU that you have at your disposal.

1) Buy orders: it is frozen the liquidity necessary to purchase the corresponding assets.

2) Selling order: it is frozen the liquidity necessary to cover the short position in the worst case scenario, i.e. a repurchase at the maximum price (100 MU). The short position can also be covered for free if the urn selected is A or B. However, considering the worst case scenario avoids bankruptcy (i.e. ending up with negative liquidity). Note that with the short sale you receive at least the limit price of your order, and therefore only the difference between 100 and your limit selling price is frozen.

It is not of crucial importance if you do not understand the details of frozen liquidity, what really matters is that you are aware that you can operate using only the available liquidity!

The trading system

The computer interface you will see during the trading activity is divided in 4 areas. From top to bottom:

a) "Information area", which contains information on:

- which of the 12 periods is being played, and the time left to insert orders;
- your total liquidity, divided between available for further exchanges and frozen.

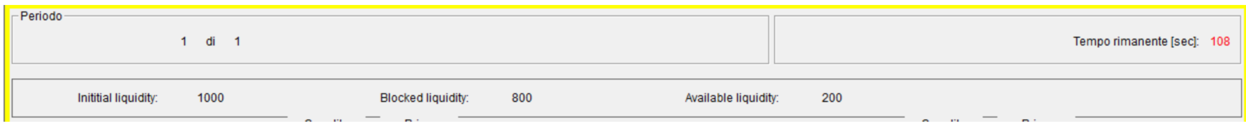

b) "Area to insert the orders": to operate on the market by inserting buy orders (on the left) and sell orders (on the right).

c) **[DA** "Book of orders": it gives an overview of the orders currently available in the market. On the left the buy orders, sorted from the most to the least convenient (the best one, i.e. that with the highest price, at the top). On the right the sell orders (the best one, i.e. that with the lowest price, at the top). For each order the quantity of assets and the price limit are shown. For example, "10 at 40MU" in the buy order column means that a total of 10 assets is ready to be purchased at a price of 40 MU each. "10 at 60MU" in the sell order column means that a total of 10 assets are ready to be sold at a price of 60 MU each. Your orders are shown in blue, those of the other traders in black.**]**

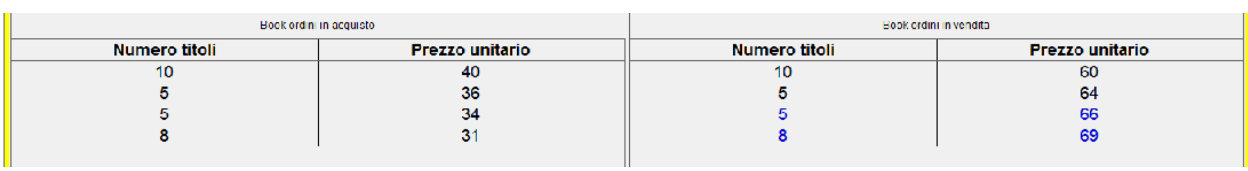

c) **[CA** "Book": it shows all your orders to buy (left) and to sell (right) the asset.**]**

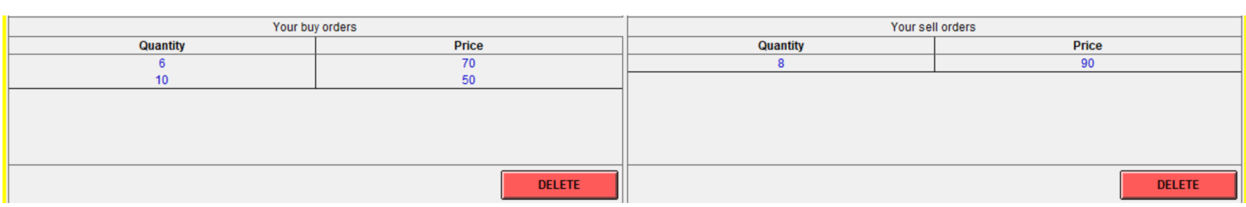

d) **[DA** ”Chronology of orders”: The left part of the screen shows all your orders, executed or pending. On the right you can instead see the chronology of the orders executed (the most recent at the top) from which you can get the trend of the market price.**]**

I miei ordini			
Tipo	**Numero titoli**	**Prezzo unitario**	**Eseguito**
Vendita	2	60	Si
Acquisto	3	40	Si
Vendita	8	60	

ELIMINA

Numero titoli	Prezzo
2	50
3	40
2	60

d) **[CA** ”Summary chart”: it displays the total number of assets that you would short sell (to the left of the vertical axes) or buy (to the right) at any price.

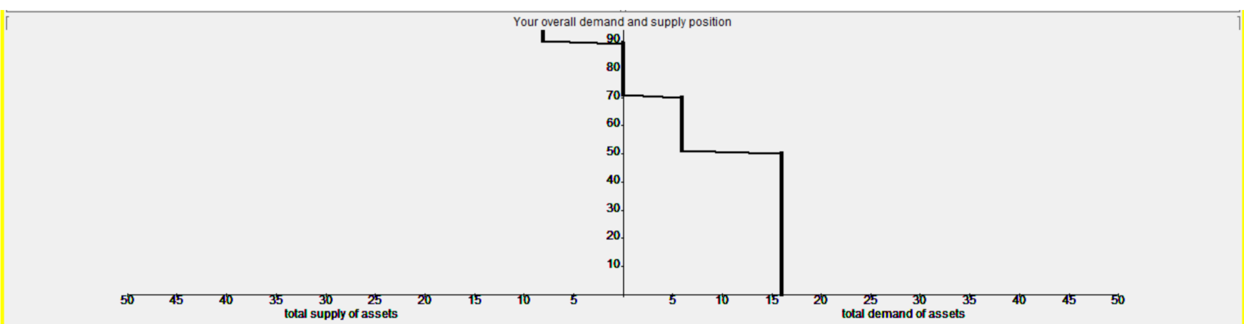

How to sum up the orders

Each order has a limit price but is also executed for ”better” prices.

For instance, looking at the buy side of the Book above (left part) we see that for a price up to 50 both orders are executed and you buy the sum of the two quantities, i.e. 16. If the price is higher than 50 the second order is not executed (because you are willing to pay up to 50 in this case), so the quantity purchased is only 6, that of the first order. This situation stays unchanged for all prices up to 50 and 70. If the price is higher than 70 it exceeds your willingness to pay even for the first order, so you do not buy anything. The mechanism is similar moving to the selling side of the book above (right part). We have already seen that for a price higher than 70 you do not buy anything, but you do not even sell until the price stays below 90. For a price of at least 90, the sell order of 8 assets is executed.

Let's make another example with the following orders:

Tuoi ordini in acquisto		Tuoi ordini in vendita	
Numero titoli	Prezzo unitario	Numero titoli	Prezzo unitario
8	10	6	30
		10	50
	ELIMINA		ELIMINA

In this case you buy 8 units at any price lower than or equal to 10, while you don't buy anything at higher prices. Moving to the selling side, for prices below 30 you do not sell anything (therefore between 11 and 29 you do not buy or sell). For a price of at least 30, the first sell order of 6 assets is executed. The situation remains unchanged for all the prices between 30 and 49. When the price is at least 50, the second order is also executed and you sell short the sum of the two quantities, i.e. 16. The Summary Chart in this case is the following:

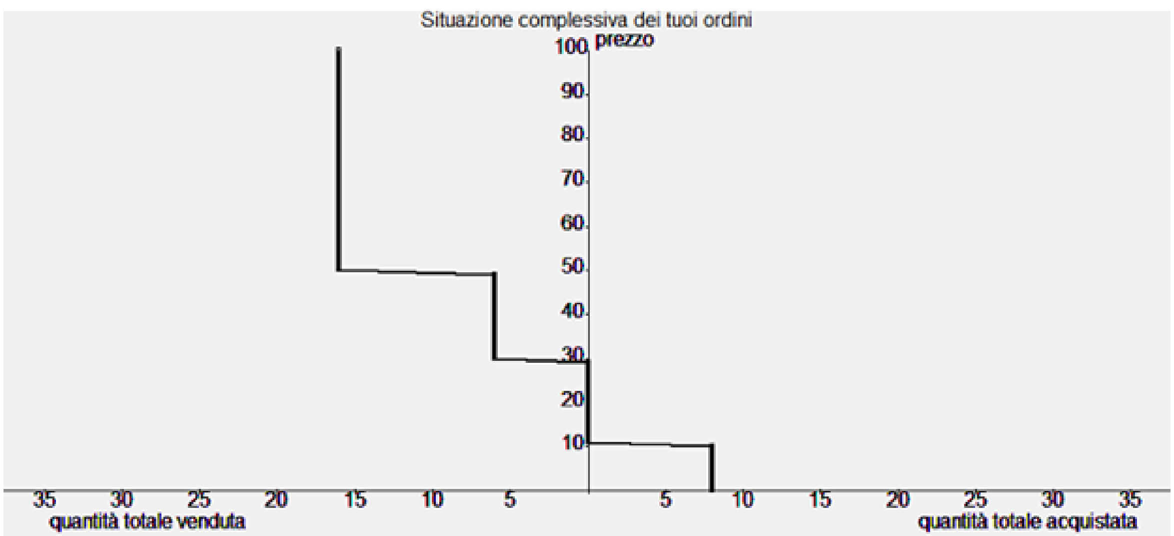

As you can see, the two graphs above differ sharply. In the first case you are willing to buy even at relatively high prices, in the second case you are willing to sell even at relatively low prices. Try to think what can determine this difference.]

How to insert and delete an order

You can insert an order in the corresponding area: purchases on the left, sales on the right. Enter the number of assets you want to exchange together with their limit price. By pressing the "Confirm" button you submit the order to the system. Multiple buy and sell orders can be inserted, provided that the necessary liquidity is available.
When inserting an order it is possible to receive the error message: "Insufficient Liquidity." The message appears when the amount of the transaction (to be spent in case of a purchase, to be kept as guarantee in the case of a short sale) exceeds the available liquidity.
Therefore, check that you did not run out of available liquidity before submitting an order. When the available liquidity is insufficient to carry out further operations you need to delete some pending order that is freezing liquidity. To do so, click with the mouse on the order in the Book, and then press the "Delete" button.
Note that if you submit a sell order with a lower limit price than a buy order of yours, the system will delete them automatically, leaving a possible residual in the Book.

Now you will see on your PC the same trading interface that you will use in the market. We ask you to do the following sequence of operations:

1. Insert a buy order of 10 assets at a price of 30MU
2. Insert a sell order of 15 assets at a price of 70MU
3. Delete one of the two orders at your choice from the Book
4. When you've done all three of them, press CONFIRM.

[QUIZ 5]

[CA only: You have now 2 minutes to practice with the same trading interface you will use in the market. You can insert buy and sell orders and see how the available liquidity, the Book of orders, and the Summary chart change accordingly.**]**

[PRACTICE PERIOD OF 2 MIN]

How the market price is determined [CA only]

a) The trading system sums up all the buy orders in a market, computing how many assets would be purchased at any price between 0 and 100. For each price this number is obtained by adding all the buy orders characterized by a higher or equal limit price. For instance, if you enter a purchase order of 10 assets at 50 MU, these assets enter the quantity demanded for all prices between 0 and 50. As 50 is the maximum you are willing to pay, this order does not contribute to the demand for prices higher than 50.

b) The trading system sums up all the sell orders in a market, computing how many assets would be sold at any price between 0 and 100. For each price, this number is obtained by adding all the sell orders characterized by a lower or equal limit price. For instance, if you enter a sell order of 10 assets at 50 MU, these assets enter the quantity supplied for all prices between 50 and 100. As 50 is the minimum you are willing to receive, this order does not contribute to the supply for prices lower than 50.

c) The trading system then compares the quantities to be sold and purchased and identifies the market price at which the number of assets bought and sold is the same, so that the exchanges can actually take place. This price maximizes the amount of assets exchanged.

Below you see two different examples of market prices. Note that at a price of 50 in the example on the right the amount demanded is greater than the one supplied and therefore the market price is greater than 50, while in the example on the left the opposite occurs. Try to imagine why this is the case.

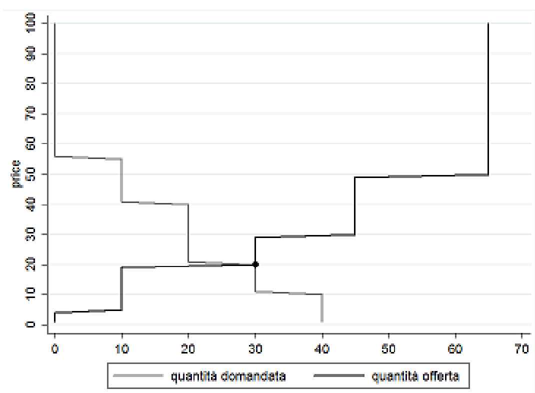

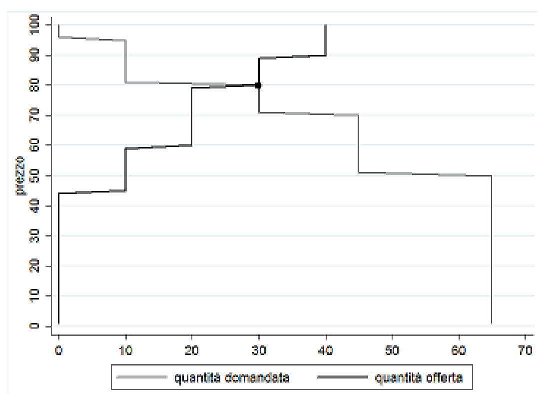

If the quantity demanded and supplied coincides in a range of prices rather than for a single price, the market price will be the average of that range. For example, if demand and supply coincides between 40 and 60, the market price will be 50.

[End of CA only]

Order execution [DA]

Two orders will be executed against each other simultaneously, only if the buy order has the same or a higher price limit as the corresponding sell order. Note that if two orders can be executed against each other, the transaction will always take place at the price limit of the previously requested buy or sell order (time priority).

Buying

Please look at the following book of pending orders, containing a list of sell orders that have been inserted into the system but not yet executed:

Book ordini in acquisto		Book ordini in vendita	
Numero titoli	Prezzo unitario	Numero titoli	Prezzo unitario
		1	50
		1	56
		1	60
		2	65

Example 1: Starting from the initial situation (0 assets, 1000 MU), assume you want to buy assets and therefore you enter a buy order "6 at 80MU" into the trading system, meaning that you are willing to buy up to 6 assets at a maximum price of 80 MU each. You see in the Order Book that one asset is offered at a price of 50 MU, 1 at 56, 1 at 60, and 2 at 65. Since you are willing to pay up to 80 MU, you will buy all 5 assets.

Important: The initial order of 6 has been split in order to be executed as much as possible (5 units), while an order for one asset remains in the trading system and will appear in the buy column of the Orders Book (1 at 80 MU). The price shown in the section "Your Orders" is your limit price for the pending order and the exchange price for the executed orders.

Following this operation in the market your account changes because you now have 5 assets in your portfolio. The amount of liquidity available is now 624 MU, because:

Money initially available	1000
Money spent to buy the 5 assets (50+56+60+2*65=296)	-296
Total liquidity	704
Liquidity blocked to ensure the execution of the pending order	-80
Available liquidity	624

Example 2: Please look at the same Order Book as in Example 1. Starting from the same initial situation (0 assets, 1000 MU), suppose you want to buy assets, but this time the maximum price you are willing to pay to buy 6 securities is 60, so you insert a buy order of "6 at 60MU." In this case you buy only 3 assets, since your limit price is not high enough to match the sell order "2 at 65MU."

Important: The initial order of 6 has been split in order to be executed as much as possible (3 units), while an order for 3 assets remains in the trading system and will appear in the buy column of the

Orders Book (3 at 60 MU). The price shown in the section ”Your Orders” is your limit price for the pending order and the exchange price for the executed orders.

Following this operation in the market your account changes because you now have 3 assets in your portfolio. The amount of liquidity available is now 654 MU, because:

Money initially available	1000
Money spent to buy the 5 assets (50+56+60=166)	-166
Total liquidity	834
Liquidity blocked to ensure the execution of the pending order	-180
Available liquidity	654

Selling

Please look at the following book of pending orders, containing a list of buy orders that have been inserted into the system but not yet executed:

Book ordini in acquisto		Book ordini in vendita	
Numero titoli	**Prezzo unitario**	**Numero titoli**	**Prezzo unitario**
1	50		
1	44		
1	40		
2	35		

Example 3: Starting from the initial situation (0 assets, 1000 MU), assume you would like to sell assets and enter a sell order “6 at 20MU” into the trading system, meaning that you are willing to sell up to 6 assets at a minimum price of 20 MU each. You do not hold this assets in your account, but you can sell them nevertheless (short selling). In other words imagine that you borrow this assets knowing that during the exchange period or at the end you need to repurchase them. You see in the book of orders that 1 asset is demanded at a price of 50 MU, 1 at 44, 1 at 40, and 2 at 35. Since you are willing to receive at least 20 MU per asset, you will sell short a quantity of 5 assets overall.

Important: The initial order of 6 has been split in order to be executed as much as possible (5 units), while an order for 1 assets remains in the trading system and will appear in the sell column of the Order Book (1 at 20 MU). The price shown in the section ”Your Orders” is your limit price for the pending order and the exchange price for the executed orders.

Following this operation in the market your account changes because you now have a number of assets equal to -5. The amount of liquidity available is now 624 MU, because:

Money initially available	1000
Money received selling short the 5 assets (50+44+40+2*35=204)	204
Total liquidity	1204
Liquidity blocked (repurchase of orders executed, worst case scenario)	-500
Liquidity blocked (repurchase of pending order, worst case scenario)	-80
Available liquidity	624

N.B. for the pending order, the system considers that if executed you will receive 20 MU, so it only takes 80 MU to cover the position in the worst case scenario.

Example 4: Please look at the same Order Book as in Example 3. Starting from the same initial situation (0 assets, 1000 MU), assume that in this case you want to sell assets but the minimum amount you are willing to receive is 40, threfore you insert a sell order "6 at 40MU." Not holding these assets in your portfolio you sell them short. In this case you sell 3 assets, since your limit price is not low enough to match the buy order "2 at 35MU."

Important: The initial order of 6 has been split in order to be executed as much as possible (3 units), while an order for 3 assets remains in the trading system and will appear in the sell column of the Order Book (3 at 40 MU). The price shown in the section "Your Orders" is your limit price for the pending order and the exchange price for the executed orders.

Following this operation in the market your account changes because you now have a number of assets equal to -3. The amount of liquidity available is now 654 MU, because:

Money initially available	1000
Money received selling short the 5 assets (50+44+40=134)	134
Total liquidity	1134
Liquidity blocked (repurchase of orders executed, worst case scenario)	-300
Liquidity blocked (repurchase of pending order, worst case scenario)	-180
Available liquidity	654

N.B. for the pending order, the system considers that if executed you will receive 3*40=120 MU, so it only takes 180 MU to cover the position in the worst case scenario.

The maximum number of assets you can buy is equal to the money available when placing the order divided by the limit price you are willing to pay. For instance, suppose that you have 1000 MU available:

- if you are willing to buy at 25: you can buy up to 40 shares.

- if you are willing to buy at 50: you can buy up to 20 shares.
- if you are willing to buy at 75: you can buy up to 13 shares.

The maximum number of assets you can sell short is equal to the money available when placing the order divided by 100 minus the limit price you are willing to receive. The reason is that for every share sold, you receive the price but you may have to spend up to 100 to repurchase the asset in the worst case scenario. For instance, suppose that you have 1000 MU available:

- if you are willing to sell at 25: you can sell short up to 13 shares (suppose you want to sell 14 shares: you would receive 14*25=350 MU, so you would have 1350 MU available that are not enough to repurchase 14 shares in the worst case scenario in which the price is 100. In contrast, if you sell short 13 shares you would receive 13*25=325 MU, that together with the initial 1000 MU would be enough to repurchase them at the maximum price).
- if you are willing to sell at 50: you can sell short up to 20 shares.
- if you are willing to sell at 75: you can sell short up to 40 shares.

Note that also in this case the situation is symmetrical between buy and sell orders.

Order execution [CA]

Buy orders are executed when the trader is willing to pay at least the market price. Buy orders with lower limit prices are not executed, those who inserted them does not receive any asset and the corresponding liquidity frozen is credited back to their account.
Following the execution of a buy order, the subject receives the asset(s) while the corresponding liquidity (market price multiplied by the amount exchanged) is withdrawn from his account.
Example: Starting from the initial situation (0 assets, 1000 MU), you want to buy assets and the maximum price you are willing to pay to buy 6 assets is 60, so you enter a buy order of "6 at 60 MU." Suppose the market price is 50. In this case you buy 6 assets paying a total of 300 MU because you pay each asset 50, not 60.

Sell orders are executed when the trader is willing to receive at most the market price. Sell orders with a higher limit price are not executed, those who inserted them does not sell any asset and the corresponding liquidity frozen to guarantee the repurchase is credited back to their account.

Following the execution of a sell order, the corresponding liquidity is credited to the account (market price multiplied by the quantity exchanged) and the assets sold short appear in their portfolio with a negative sign.
Example: Starting from the initial situation (0 assets, 1000 MU) you want to sell assets and the minimum price you are willing to receive to sell 6 assets is 40. Not holding these assets you sell them short. Suppose the market price is 50. In this case you short-sell 6 assets, therefore holding a negative balance (-6) and you receive 300 MU because you receive 50, not 40 for each asset.

[ONLY IN CA: Partial order execution

If the quantities demanded and supplied do not exactly match at the market price, it is possible that some orders are not executed at all or in part.
For example, if at a market price of 50 the quantity demanded is 60 but the quantity supplied is 65 it is possible to exchange 60 assets at most. Sell orders for 5 assets will not be executed because there is no counterpart willing to buy them at that price. Likewise, if the quantity demanded is 65 and the quantity supplied is 60, buy orders for 5 assets will remain unexecuted.
Priority is given to the execution of buy orders with the highest limit price and to sell orders with the lowest price limit. In case of a tie, priority is given to the order inserted first. **]**

[ONLY IN CA: Outcome of the trading period

At the end of each trading period you will see a screen that summarizes:

1. The market price in that period
2. Your account, including the number of assets (purchased at the market price if the balance is positive, or sold short if the balance is negative) and your liquidity.**]**

Note that you will know which urn was used in each period only at the end of the whole experiment when your earnings will be determined.

4. What are your earnings in the market?

At the end of the experiment one of the 12 trading periods will be randomly selected and used to determine your compensation. Earnings are given by the sum of
- total liquidity at the end of the trading period;
- value of your portfolio of assets: 0 if the urn selected is A or B; 100 MU multiplied by the number (positive or negative) of assets if the urn selected is C or D. In the case of a net short position

(negative balance), the value of the portfolio is equivalent to an automatic repurchase at the final value of the security: 0 (urn A or B) or 100 (urn C or D).

Summary of the procedures

We are now going to start the trading phase, which consists of a total of 12 periods. In each period you first receive the signal about the number of blue marbles, after which you will be asked to estimate the probability of each urn. Then the 3 minute **[CA:** 2 minute **]** period in which you can insert the orders will start. At the end of the trading period you will again be asked to estimate the probability of each urn.

After the 12 trading periods, we will proceed:

1. Drawing the outcome of the Investment Game (Head or Tail);
2. Drawing the relevant period (from 1 to 12) relevant for the earnings of your market activity;
3. Drawing the period and the phase relevant for the estimation of the probabilities of the urns (from 1 to 12 but different from that relevant for the earnings in the market), and of the number between 0 and 100 that is used to compare the accuracy of your estimates.
4. Finally, we will ask you to fill out a quick questionnaire.

Summary of your earnings in the experiment

Your earnings in the experiment are the sum of the payoffs obtained in the various phases:

1. Investment Game;
2. Estimate of the probability of the urns;
3. Market activity.

This sum divided by 100 represents your payment in euro, to which a show up fee of 2.5€is added. The total amount will be paid to you anonymously at the end of the experiment.