Risk Aversion and Information Aggregation in Asset Markets ☆

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

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

PRELIMINARY DRAFT, PLEASE DO NOT CITE.
COMMENTS WELCOME.

Abstract

The paper investigates the relation between the risk preferences of traders and the information-aggregation properties of an experimental asset market. In the experiment, we manipulate the level of risk aversion of the traders across markets. We show that market-clearing prices should be more informative when traders are more risk-averse. The evidence is not consistent with this prediction and is close to the risk-neutral benchmark. However, individuals are shown to be risk averse both in the market and in an independent task. We show that this conflict descends from individuals expressing risk-aversion in two different ways. The first is the classical curvature of the demand schedule; the second is a attitude to act as if one had less information than she actually has. While being suboptimal, this attitude is a way to be cautious on the market. Indeed, despite playing in a complex asset market, subjects' individual behavior clearly retains the footprint of their risk attitudes as measured through an independent task.

JEL Classifications: C81; C91; D81

Contact: `antonio.filippin@unimi.it` (Antonio Filippin), `marco.mantovani@unimib.it` (Marco Mantovani)

November 23, 2018

1. Introduction

The paper investigates the relation between the risk preferences of traders and the information-aggregation properties of an experimental asset market.

Asset markets have a central role in aggregating dispersed information, and their performance on this dimension is the main yardstick to evaluate market efficiency. Indeed, estimating traders' beliefs from prices is a common empirical exercise in financial markets in general (e.g. Alti and Tetlock, 2014; Cipriani and Guarino, 2014; Easley et al., 1997). Some financial markets, called prediction markets, exist with the unique objective of allowing for inferences from prices to beliefs (e.g. Arrow et al., 2008; Wolfers and Zitzewitz, 2004). Therefore, they rest on the validity of those inferences, which is based on restrictive assumptions about the risk preferences of the traders (Manski, 2006). While trading a risky asset is an archetypical instance of decision under uncertainty, the empirical evidence on how risk preferences affect information aggregation in markets is scant.

Our experimental exercise is characterized by two main distinctive features. First, our market is a call auction that allows to directly observe individual demand schedules. Second, we elicit individuals' risk preferences through an independent task, and use this measure to manipulate risk-aversion at the market level. That is, we form markets with markedly different risk-aversion levels. Across these markets, the level of risk aversion should have an unambiguous effect on market prices: prices should be more informative (about the state) in more risk-averse markets.

We show that traders' risk preferences are reflected in market activity (orders, volumes), but do not affect market prices in the way predicted by the theory. Thanks to the possibility to analyze individual demand schedules, we identify a novel and suboptimal way in which risk preferences are expressed in the market. that explains the previous findings. On top of submitting more curved demand schedules – as expected under expected utility maximization –, more risk-averse subjects reduce their market exposure by acting *as if* they had less information than they do. While the classical 'curvature' effect of risk aversion would lead to more informative prices in more risk-averse markets, this second effect – which is an 'intercept' effect – goes in the opposite direction, and the two offset each other. As a corollary, our results suggest caution in inferring risk preferences from market prices. On the other hand, our results are positive on the value of risk-elicitation methods as tools to capture stable features of individual preferences: despite playing in a complex asset market, subjects' individual behavior clearly retains the footprint of their risk attitudes as measured through an independent task. This footprint is all the more visible once we embrace the broader class of cautious market behavior that we have identified.

In the experiment, subjects trade an Arrow-Debreu security in a two-state economy. Subjects' hold a common prior about the state and receive information in the form of imprecise signals. Signals induce heterogeneous posterior beliefs. Since failures in Bayesian updating may influence prices, we elicit beliefs in an incentive compatible way. Given the state, subjects hold homogeneous induced preferences over the asset, so that trades are zero-sum. Therefore, trading is only driven by the different information held by the traders. All traders are endowed with the same amount of a numeraire good and have no holding in assets: they can only sell through short sales. This allows not to deal with the joint distribution of information and endowments and avoids constraints on traded volumes. Net short positions are covered at the actual value of the asset once the state is revealed.

The market is a single call auction. Subjects have two minutes to place limit bid and ask orders in isolation. The combination of these orders form their demand schedule – a quantity supplied or demanded for each possible price. At the end of the two minutes, the aggregate demand schedule identifies a market-clearing price, and trading occurs at that price.

Prior to entering the market, we elicit traders' risk preferences using the Investment Game Gneezy and Potters (1997). We manipulate risk preferences at the market level by grouping traders according to their choices in the Investment Game. In each session, we form two markets, one formed by the eleven subjects that are the most risk averse, the other by the eleven subjects that are the least risk averse.

We generate hypotheses studying the so-called prior information equilibria of the asset market. In this model, traders do not behave strategically: they bring the information they have to the market through their demand schedules, but do not internalize the information contained in market-clearing prices. In other words, when deciding how much to buy at a certain price, they do not anticipate the beliefs they will have if that price happened to be the market-clearing price.

In our set-up, the prior information equilibria provide clear-cut predictions on the role of risk aversion in information aggregation: the more risk averse are the traders, the more market-clearing prices are informative about the state.

The hypothesis on the role of risk-aversion on prices hinges on the prior information model, which posits that traders do not anticipate the informative content of prices when submitting their orders. Because the effect of this type of strategic behavior could be confounded with the effect of risk aversion, we also run control sessions under a non-strategic institution where prices cannot depend on traders' information as their are drawn randomly from a known distribution. This random-price mechanism – similar in spirit to Becker et al. (1964) – is used to test for the validity of the prior information model, and the prior information model passes the test.

Results can be summarized as follows. The informational content of prices is not affected by traders' risk aversion. Prices are informative, in the sense that they respond to the information that is given to traders. However, they are close to risk-neutral ones – i.e., minimally informative relative to any other level of risk-aversion. Elicited beliefs show this is not due to failures in belief updating.

Traders are risk-averse according to their choices in the Investment Game. These choices correctly predict trading volumes in the market: more risk-averse individuals (and markets) trade significantly lower volumes. What is the mechanism through which risk-averse traders lead to risk-neutral prices?

To answer, we analyze individual demands schedules. These are shown to be at odds with risk neutrality. Indeed, risk-aversion coefficients estimated from the demand schedules are larger than those estimated from the Investment game. Demand schedules show individuals act as if they had less information than they actually have, both in theory and according to their elicited beliefs. While estimated risk aversion parameters are able to capture their actual curvature, they fail at capturing the position of demand schedules – i.e. an intercept effect. We then re-estimate the individual risk-aversion coefficient and a second parameter capturing this attitude. We show that acting as if one had less information reduces exposure to the market and represent a way to trade lower expected earnings with a lower variance of earnings. Thus, while individuals have a more convenient way of doing it (by submitting a more curved demand schedule), this behavior represents a suboptimal way of expressing risk aversion in the market. Consistently, we find that both parameters are significantly correlated with choices in the risk elicitation task. These two ways of expressing risk aversion in markets drive market-clearing prices in opposite directions, which explains both why we did not find differences in the informational content of prices across markets with different levels of risk aversion, and why prices are, on aggregate less informative than what elicited risk-aversion levels would imply.

Our results provide a leap forward in the understanding of risk aversion as commonly measured in the lab. We show that elicited risk aversion may fail to represent a stable and portable construct, contrary to what the curvature of the utility function represents, because they possibly encompasses additional behavioral characteristics. According to our findings, risk attitudes constitute a wider construct that captures a more general reluctance to act, incorporating for instance the unwillingness to exploit the information available.

2. The theoretical role of risk aversion

There are N traders facing uncertainty regarding two ex-ante equally likely events, $e \in \{R, B\}$. An Arrow-Debreu security is traded on a market. The security pays 100 to its owner if $e = B$, and pays 0 is $e = R$.[1]

[1] For the sake of consistency with the experimental design we use $p = 100$ and express probabilities and beliefs in percentage points.

Traders' preferences are represented by an individual utility function over wealth levels, featuring *constant relative risk aversion* (CRRA). Risk preferences are allowed to differ across traders. That is:

$$u_i(w_i) = \begin{cases} \frac{w_i^{1-\theta_i}}{1-\theta_i} & \text{if } \theta_i \neq 1 \\ ln(w_i) & \text{if } \theta_i = 1 \end{cases} \tag{1}$$

Each trader forms a belief $b_i \in [0, 100]$ exploiting his private information, where b_i represents his subjective probability that $e = B$. In Section 3 we clarify that we do not impose that subjects use their private information in a Bayesian manner, but for the purpose of this Section we simply take b_i as given.

Traders enter the market with an equal endowment m of a numeraire good, whose value is state-independent. One unit of the numeraire good pays one unit of wealth, in either state. Given this assumption we will simply refer to the numeraire good as 'monetary endowment' for the sake of simplicity. Since there are no endowments of the security, sales can occur only through short selling. Short positions are closed at the end of the market at the actual value of the security, given the realized state. That is, after the closure of the market, sellers buy the asset for 100 if $e = B$, for free if $e = R$.[2]

Under a Single Call Auction (CA) traders submit a demand schedule $q_i(p)$, for $p \in [0, 1]$, where a negative demand indicate a short position at a given price. Demands must satisfy a no-bankruptcy condition: traders' obligations cannot exceed their monetary endowment, independent of the actual state. The market mechanism aggregates individual demands and trades are executed, according to the individual demands, at a unique market-clearing price. That is, a price p^* such $\sum_i q_i(p^*) = 0$.[3]

We assume that traders behave as expected utility maximizers and act as price-takers on the market. Furthermore, we assume here the so-called *prior information* model. This model posits that traders submit demand schedules according to their beliefs and preferences, but disregard the informational content of prices. In other words, they do not conjecture what distribution of others' beliefs (and, thus, information about the state) would sustain a certain market-clearing price. Under this assumption traders provide information without extracting information from the market before the closure, and prices simply aggregate the information the traders have *prior* to entering the market. As explained in Section 3 we test the validity of the prior information assumption in a specific treatment, where we employ

[2]Given these rules, this set-up is isomorphic to a two-states/two-assets environment: holding a short position in our set-up is identical to holding a long position for a security that pays only when $e = R$. Considering a single security simplifies the experimental task for subjects.

[3]Let us assume for the moment that the demand schedules behave so that such a price exists and is unique. In Section 3 we will deal with the possibility of imperfect market clearing and multiple market-clearing prices.

a price formation mechanism á la Becker et al. (1964) in which prices are exogenous with respect to the demand schedules by construction.

As shown in for instance in Gjerstad (2005), solving for the trader's maximization problem under these assumptions yields:

$$q_i^*(p, b_i, \theta_i > 0) = \frac{(1-p)^{\frac{1}{\theta_i}} b_i^{\frac{1}{\theta_i}} - p^{\frac{1}{\theta_i}} (1-b_i)^{\frac{1}{\theta_i}}}{(1-p)\, p^{\frac{1}{\theta_i}} (1-b_i)^{\frac{1}{\theta_i}} + p\,(1-p)^{\frac{1}{\theta_i}} b_i^{\frac{1}{\theta_i}}} m, \tag{2}$$

for the case of risk averse players, and:

$$q_i^*(p, b_i, \theta_i \leq 0) = \begin{cases} \frac{m}{p} & \text{if } b_i > p \\ \left[-\frac{m}{1-p}, \frac{m}{p}\right] & \text{if } b_i = p \\ -\frac{m}{1-p} & \text{if } b_i < p, \end{cases} \tag{3}$$

for the case of risk neutral and risk loving traders.

Risk aversion plays a crucial role in shaping the market outcomes both at the individual and consequently at the aggregate level. The optimal demand schedules for different levels of the CRRA coefficient are depicted in Figure 1. The demand of risk neutral and risk loving players reacts sharply to changes in the prices around their belief. These traders invest their entire endowment in buying or selling the security, depending on whether the price is (even slightly) below or above his belief, respectively.

The demand of risk averse traders responds instead smoothly, with the amount invested that correlates with the distance between the price and the trader's belief. The observation that at any given price the amount invested decreases with the degree of risk aversion allows us to identify a testable implication at the individual level.

Hypothesis 1: The degree of risk aversion negatively correlates with the trading volumes at the individual level.

Different degrees of risk aversion also have implications at the aggregate level because the different individual demands shape the market price. For the sake of simplicity we assume here that the traders are characterized by the same degree of risk aversion.[4] Approaching risk neutrality ($\theta \to 0$ for all i) what happens is that $p^* \to 50$, i.e. the market clearing price tends to the prior probability of the event. Mild degrees of risk aversion ($0 < \theta < 1$) replicate the so-called favorite-longshot bias, an empirical regularity according to which unlikely states are over-priced, and likely states are under-priced: $p^* > \bar{b}$ when $\bar{b} < 50$, and

[4]Conditions under which similar results hold with heterogeneous risk preferences are shown by He and Treich (2013) theoretically, and by Fountain and Harrison (2011) with a simulation exercise. In Section 5 we extend the empirical analysis accounting for the individual heterogeneity in our data.

Figure 1: Optimal demand shcedule and risk aversion

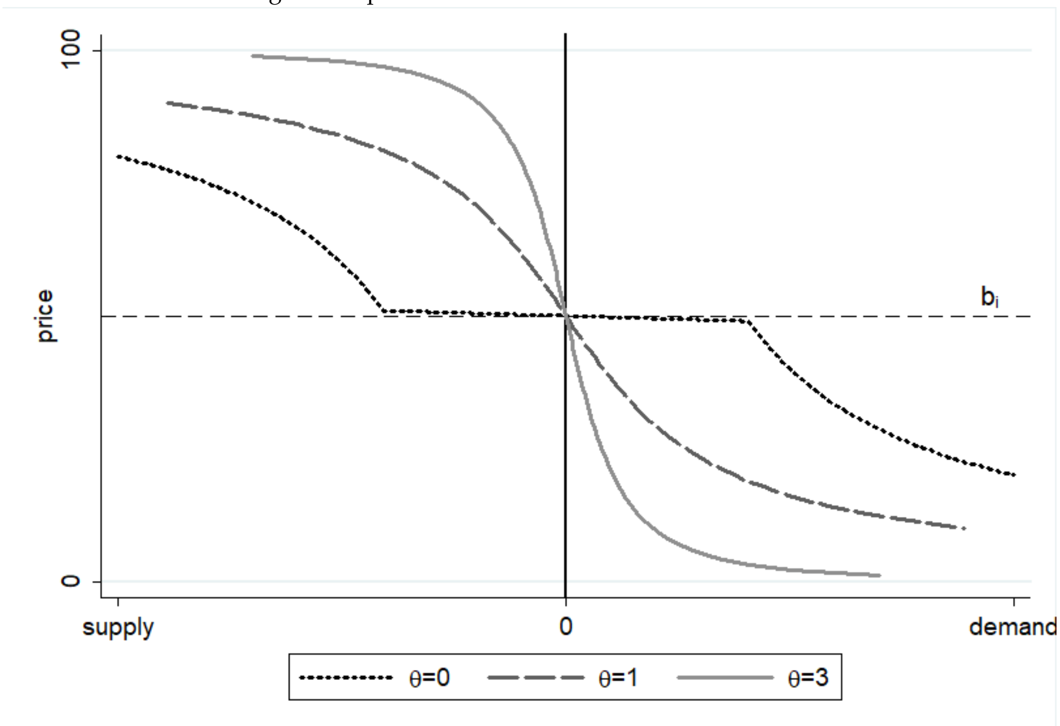

$p^* < \bar{b}$ when $\bar{b} > 50$ (Manski, 2006). Log-utility ($\theta = 1$) sets an important benchmark, as in this case the market-clearing price is equal to the average belief of traders, $p^* = \bar{b}$, defining a proper prediction market (Wolfers and Zitzewitz, 2006). As risk aversion further increases ($\theta > 1$) the market price exceeds the average beliefs: $p^* > \bar{b} > 50$ and $p^* < \bar{b} < 50$. Note that as long as the market as a whole has some information to share, i.e. when $\bar{b} < 50|e = R$ and $\bar{b} > 50|e = B$, risk aversion has direct consequences on the informational content of the prices. Given that theoretical predictions are always doomed to fail when stretched to hold too precisely, we can summarize the predictions above identifying a more conservative testable implication at the aggregate level:

Hypothesis 2: The distance between market prices and average beliefs is minimized around log utility.

This Hypothesis will be defined more precisely in Section 5 incorporating the parametrization of our experimental design.

2.1. Informative Prices

We now analyse the effect on individual demands if we relax the prior information assumption. In this case, traders take into account the information conveyed by any conceivable market-clearing price about the state of nature when submitting their limit orders. In the limit case in which all the relevant information is contained in the market price, traders cannot profit from their private information. At any price they should be neither buyers nor sellers ($b_i = p$ in Eq. 2), and the straightforward outcome is therefore no trade (Milgrom and Stokey, 1982). There are several reasons why such a scenario is unrealistic. For instance, if the number of traders is small or in presence of noise traders the private information retains a positive value and contribute to shape the individual behavior together with the prices.

Figure 2 illustrates what happens to the individual demand in our setting when traders use the market prices to complement their private information. The dashed line represent the benchmark demand for a subject characterized by $b_i = 50$ and $\theta_i = 1$ in the private information model. At $p = 50$ the information extracted from the market concur with the private information, so that the choice will not change (point A). Consider now a higher price, e.g. $p = 70$. The trader should internalize that the market is signalling a probability that the asset will pay 100 higher than he believes, and readjust his beliefs accordingly. Let's assume for the sake of simplicity that the subject weights linearly the two sources of information:

$$b'_i = \alpha p + (1 - \alpha) b_i. \tag{4}$$

The relative weight assigned to his private information and to the market price will determine his choice, with prior information ($\alpha = 0$) and no trade ($\alpha = 1$) being the two extremes. How much information the trader extracts from the market depends on α. For instance, if $\alpha = 0.5$ the trader will behave as if he now believes that the probability that $e = Blue$ is equal to 60 when considering $p = 70$. As a result he behaves as if his individual demand was the dotted line q_{60}, with point B_2 representing his net demand at $p = 70$ as compared to B_1 in the prior information case. A similar procedure is illustrated for $p = 30$, leading to a net demand corresponding to point C_2. The envelope of the results obtained iterating the mechanism above for any price $p \in (1, 100)$ represents the individual demand of a trader characterized by $\alpha = 0.5$ (solid line in the graph).

The informative content of the price has the effect of increasing the curvature of the individual demand: the larger is α, the more curved is the demand. Note that extracting information from the market is confounded with a higher degree of risk aversion. Given that the focus of our paper is on the role played by risk aversion, the prior information assumption turns out to be crucial in our setting. For this reason, in Section 3 we test the validity of the prior information assumption using a specific treatment, where we employ a price formation mechanism á la Becker et al. (1964) in which prices are exogenous with respect to the demand schedules by construction.

Figure 2: Optimal demand schedule with informative prices

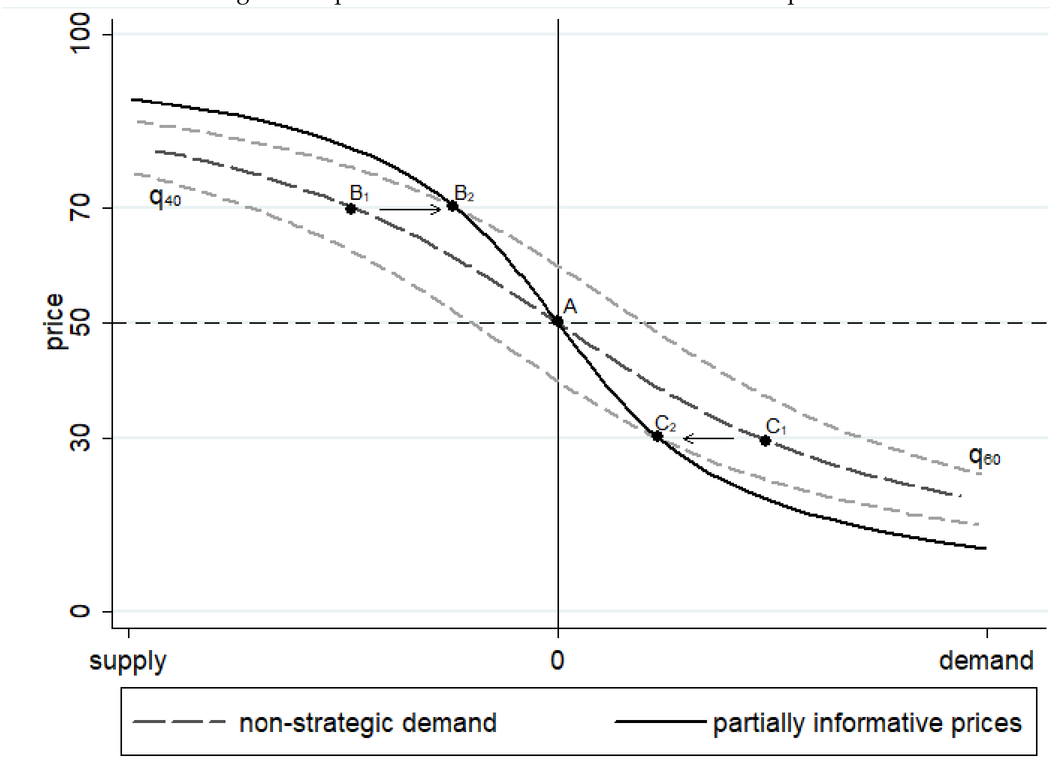

3. Design

Elicitation of risk preferences. The purpose of the experiment is to evaluate the relation between the distribution of risk preferences in the population and the properties of market prices in terms of information aggregation. At the beginning of the experiment, we elicit subjects' risk preferences using the Investment Game (Gneezy and Potters, 1997). In this task, subjects have to decide how to allocate a given endowment of 200 MU between a safe account and a risky investment. The latter yields 2.5 times the amount invested or zero, with equal probability. We chose this risk elicitation task because, as reported by Crosetto and Filippin (2016), it is more suitable to scan risk preferences around logarithmic utility than, for instance, the Holt and Laury task (Holt and Laury, 2002) or the Bomb Risk Elicitation Task (Crosetto and Filippin, 2013). As discussed in Section 2, this level of risk aversion is crucial in our framework, because it is the only one that guarantees convergence of the market-clearing price to the average belief of the traders (Wolfers and Zitzewitz, 2006).

We also divide each session in two groups of 11 traders according the their choice in the Investment Game, separating the subjects above and below the median. By doing so we exogenously induce variability in the distribution of (elicited) risk preferences across markets, while at the same time minimizing the heterogeneity within market.

The asset market. 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. The 11 traders are not informed of which urn has been selected, but know that each urn is selected with equal probability. That is, all subjects start with the same (common knowledge) prior distribution of outcomes.

A simple asset called "Majority Blue" is traded in the market. If the urn is C or D, the event "the majority of marbles are Blue" occurs ($e = Blue$) and every asset will pay the owner 100 monetary units (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 occur ($e = Red$) and every asset will pay the owner 0 MU.[5]

Subjects 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. $s \in \{x-5, x+5\}$ where x is the true number of blue marbles. Given the urn, subjects receive each signal with the same probability. In fact, each of the 11 subjects is randomly assigned one of the 11 possible signals, as illustrated in Table 1. For instance, if the selected urn is A (47 blue), the 11 subjects will receive one of the signals ranging from 42 to 52. The procedure that generates and distribute the signals is common knowledge.

	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

Table 1: Signals

Given the signal, Bayesian updating of the prior about the event generates the posterior beliefs $p(e = Blue|s)$ reported in the last row of Table 1. Some signals ($s \leq 45$ and $s \geq 55$) fully reveal the outcome, leading to $p(e = Blue|s) \in \{0, 1\}$. Others ($s = 46, 47, 53, 54$) are partially informative, i.e. $p(e = Blue|s) \in \{\frac{1}{3}, \frac{2}{3}\}$. Finally, some signals ($s \in [48, 52]$) are uninformative and therefore $p(e = Blue|s) = \frac{1}{2}$.

The combination of the underlying urn with the structure of the signal constitutes a manipulation of the amount and the distribution of information in the market. Aggregating

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

$p(e = Blue|s)$ over subjects, the average beliefs about the state $e = Blue$ turn out to be 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. According to the theoretical model, whether market prices trace these values or not should depend on the degree of risk aversion. With logarithmic utility ($\theta = 1$) prices should reflect exactly the average belief of the population, while lower (higher) degrees of risk aversion should determine a flatter (steeper) pattern.

Belief elicitation. Even assuming log utility, prices can diverge from the pattern above because subjects do not update their beliefs in a Bayesian manner, or more generally if they misperceive the objective probabilities. It is also possible that failures of Bayesian updating correlate with risk preferences, further confounding the results on the role played by risk attitudes in the process of information aggregation. To address this issue, we ask subjects to report their subjective probability that each urn has been selected.

The elicitation of beliefs is incentivized using the Binarized Scoring Rule (BRS) (Hossain and Okui, 2013). The BSR consists in comparing the sum of squared errors of the reported beliefs (normalized between 0 to 1) with a 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. The BRS is isomorphic to the Quadratic Scoring Rule (QSR) in terms of expected reward, but instead of paying different amounts according to the accuracy of beliefs like the QSR, the BSR pays a different probability of receiving the higher of two discrete amounts. Since in the BSR the variance of the outcomes cannot be reduced, this procedure turns out to be incentive-compatible regardless of subject's risk attitudes. The optimal choice always requires to maximize the likelihood of getting the high amount and therefore to truthfully reveal one's beliefs.[6]

Trading and market institution. Subjects receive a monetary endowment of 1000 MU in each trading period. They have two minutes in which they can place bid and ask orders for the asset. Since subjects have no endowment of assets, sales are implemented via short selling. Short positions are covered at the end of the trading period at the actual value of the asset. That is, subjects who are net sellers at the market-clearing price, buy back the assets for 0 MU if the urn is A or B, for 100 MU if the urn is C or D. No-bankruptcy is ensured freezing liquidity for pending orders, making sure that the net demand does not require more than the endowment for any possible market price.

The market institution better suited to investigate the role played by risk aversion in the aggregation of information, and therefore in the determination of the equilibrium prices, is

[6] The QSR instead induces a truthful revelation of beliefs only for risk-neutral subjects, but it is not incentive-compatible in general. For instance, a risk-averse subject may prefer to smooth the reported beliefs because the reduced variance of the outcomes more than compensates the loss in terms of expected reward.

the Call Auction (henceforth CA). In our opinion the CA is better than other market institutions such as the Double Auction for at least two reasons. First, in a CA prices are formed simply aggregating the individual demands, i.e. without the interaction among traders that would blur the link between risk attitudes and equilibrium prices. Second, the CA provides individual net demands that allows us to estimate the individual degree of risk aversion, which can then be compared with the elicited ones in order to test their consistency.

As discussed in Section 2, however, subjects may extract information from the market prices, contrary to what assumed under the prior information model. This holds also in a CA, even if subjects do not observe realized prices (as in the Double Auction) while placing orders. Indeed, they may conjecture that market-clearing prices are influenced by the signals received by others. If this is the case, they could then internalize in their orders the information contained in the signals of other traders. There is experimental evidence showing that agents in a CA treat the price simply as an opportunity cost (Biais et al., 2017; Ngangoue and Weizsacker, 2015). However, conditioning one's behavior on the informative content of the price would act as a confounding factor for our main variable of interest, i.e. risk aversion. Therefore, in order to test the validity of the prior information assumption, we run two treatments in a between subject design manipulating the price formation mechanism *ceteris paribus*. The sequence of urns and distribution of signals in every session is the same in the two treatments, in order to keep the informative structure constant across conditions.

Treatment CA. In this treatment the equilibrium price at which orders are executed is computed as the price that equalize aggregate demand and supply, maximizing the volume of trades.[7]

Treatment RPM. In this condition the execution of the orders follows a Random Price Mechanism (RPM), á la Becker et al. (1964). For each subject, the price is randomly drawn from a known distribution. Since this distribution is independent of the selected urn, and this is known, prices do not have any informative content. The probability distribution of the prices mimics the empirical frequency of prices observed in the CA treatment, and is represented in Table 2.[8] Given this procedure, every subject constitutes one independent observation in this treatment.

After testing the goodness of the prior information assumption, the results in the CA can also be used to retrieve the individual coefficient of risk aversion from the net demands.

[7] In case demand and supply are equal for a range of prices the average price 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.

[8] The reason why prices are not drawn from a uniform distribution is that by observing more often prices in the tails of the support as compared to the CA treatment could induce a different behavior regardless of the informative content of the price.

Table 2: Distribution of prices in treatment RPM

Price window	1 – 14	15 – 29	30 – 44	45 – 55	56 – 70	71 – 85	86 – 99
Probability	.01	.09	.20	.40	.20	.09	.01

Notes: the table reports the distribution from which prices are extracted in treatment RPM. The procedure works as follows. First one price window is selected. Each price window is selected with the corresponding probability. Second one price within the price window is selected at random with uniform probability.

Therefore, the consistency of the elicited risk preferences and the behavior in the market can be replicated at the individual level, besides in aggregate terms (see Hypothesis 2 above):

Hypothesis 3: The individual coefficient of risk aversion estimated from the market behavior is consistent with the measure elicited with the Investment Game.

4. Procedures and Payment scheme

Sessions were run between February and September 2018 at the EELAB at the University of Milan - Bicocca. The experimental software was developed using Z-tree (Fischbacher, 2007). All sessions follow identical procedures. Upon arrival, subjects are randomly assigned to cubicles in the lab. They first face the Investment Game. The software uses their choices to assign them to one of to markets of 11 traders. If their choice in the Investment Game is below the median of the session, the subject is assigned to the High risk aversion market. Otherwise, he is assigned to the Low risk aversion market.

Subjects then receive detailed instruction on the rules and the working of the market. During the instructions, they are asked to answer a battery of quizzes that assess their comprehension of the various parts of the instructions.[9] The reading of the instructions would move on only once all subjects have cleared the quiz. For each quiz, we record the number of mistakes each subject makes before clearing the quiz.

Subjects then play 12 trading periods. Within each period, they first receive their signal and have 30 seconds to report the probability that each urn has been selected. Then they have 2 minutes to insert their limit orders. A graphic representation of their demand schedule updates in real time on their screen each time they insert or erase an order. The software checks that a no-bankruptcy condition is satisfied independent of the actual state of the world before accepting an order, and returns an error message in case the condition is not satisfied. At the end of the trading periods, subjects are informed of the price in that

[9] 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 fifth and last one, the working of the market interface, and also includes two minutes to interact freely with the interface.

period, their liquidity and asset portfolio, but not of the selected urn. Then, they are asked to to report again the probability that each urn has been selected.

At the end of the experiment the computer selects (i) by tossing a virtual coin, for each participant, the outcome of the Investment Game; (ii) at random , for all participants, one period to be used for payments of the trading task; (iii) at random, for each participant, one of the 24 measures of beliefs. The selected measure of beliefs must be in a period that is not relevant for the payment of the trading task, so as to avoid hedging strategies across tasks. To compute payments of the belief task according to the BSR, one random number is assigned to every participant. Random numbers are different for every participant, so as to avoid social comparison effects. To ensure credibility of our procedures, subjects at the end receive detailed information about the distribution of all random draws. After that, they receive feedback on their payment, and fill in a short questionnaire, after which the experiment is over.

For treatment CA, we collect data from 10 sessions, or 220 experimental subjects. These correspond to 20 independent observations – 10 High and 10 Low risk aversion markets – with data on 240 trading periods, 60 for each urn. For treatment RPM, we collect data from 38 experimental subjects/independent observations, divided into 2 sessions. Sessions lasted about 2 hours. The average payment was 16.2 €.

5. Results

We open the section with the comparison of treatments CA and RPM, in order validate the prior information assumption. We then look at the relation between elicited risk aversion and trading volumes, and test Hypothesis 1. The following subsection tests Hypothesis 2, investigating if and how the average level of risk aversion in a market affects the market-clearing price. We then use individual demand data to estimate a subject's risk aversion directly from his market activity. The comparison of these estimates with aggregate market outcomes induces us to enrich the specification of the individual demand. In particular, we broaden the possibility to display a reluctance to act beyond what descending from the curvature of the utility function by encompassing into the model a partial use of the information available. Finally, we look at how the estimates of risk aversion and reluctance to act obtained from the market activity correlate with the measure of risk aversion elicited through the investment game, thereby testing Hypothesis 3.

5.1. Test of the prior information model

The theoretical analysis of risk aversion builds on the prior information model – i.e., the assumption that subjects do not exploit any informative content of market-clearing prices. If they did, risk aversion would be confounded with strategic behavior (see Figure 2). We test the validity of the prior information assumption by comparing behavior in treatment

CA to that in treatment RPM, where prices are uninformative by construction because they do not depend on subjects' behavior.

Figure 3 shows the average demand schedule in the two treatments over all periods. If subjects anticipated the informative content of prices in treatment CA, this would result in a steeper demand schedule than in the RPM treatment. This effect is not visible in Figure 3. To perform a formal test of the prior information model, we estimate, for each subject, the slope of his demand function, aggregating over all periods. Under CRRA and prior information model, individual demand is given by equation 2, so we obtain the individual parameter with a Maximum Likelihood estimation. The null hypothesis is that individual θ_{mkt} coefficients come from the same distribution in the two treatments and is tested using one independent observation per subject performing a Mann Whitney U-test. [10] The null hypothesis of equality across treatments is not rejected ($U = -1.097$, $p = .273$) and therefore we conclude that individual behavior does not differ in treatment CA and RPM. [11] Given that the prior information model passes the test, in what follows the analysis will focus on the data of the CA treatment only, while the RPM will be disregarded.

5.2. Elicited risk aversion: Individual choices

In the Investment Game (Gneezy and Potters, 1997) subjects invest on average 93.5 out of 200 ECU, while the median choice is 100. An investment of about 50% is in the lower part of the range of results commonly observed in this task (see Charness and Gneezy, 2012, and references therein), and corresponds to an average CRRA coefficient $\theta_{inv} = 0.71$ (median 0.32). The median seems to signal a low degree of risk aversion, but this is mainly the result of the inner working of the task. In fact, the transformation of choices into θ_{inv} coefficients is highly non-linear: it is very sensitive to changes in the choice for low amounts invested, it corresponds already to a low level of risk aversion ($\theta_{inv} = 0.32$) when half of the endowment is kept, and it decreases then very slowly for more risk taking decisions (see Figure 3 in Crosetto and Filippin, 2016). Decisions in our experiment are indeed far from the implication of risk neutrality, i.e. that the whole endowment should be invested, since only 10% of the subjects opt for such a choice. The cumulative distribution of all the choices is reported in Figure A.11 of Appendix A.

[10] To avoid confusion from now on we call θ_{mkt} the risk aversion coefficient estimated from market net demands, while we refer to risk aversion coefficient elicited in the Investment Game as θ_{inv}. To avoid to overload the notation we omit the subscript i when referring to individual parameters.

[11] Note that the negative U statistic indicates that the curvatures are more often higher in the RPM than in the CA treatment, which is the opposite of what a violation of the prior information assumption would imply. The result of the test is the same with the comparison of the slope of a linear approximation of the individual demands across treatments.

Figure 3: Average demand in treatment CA and RPM.

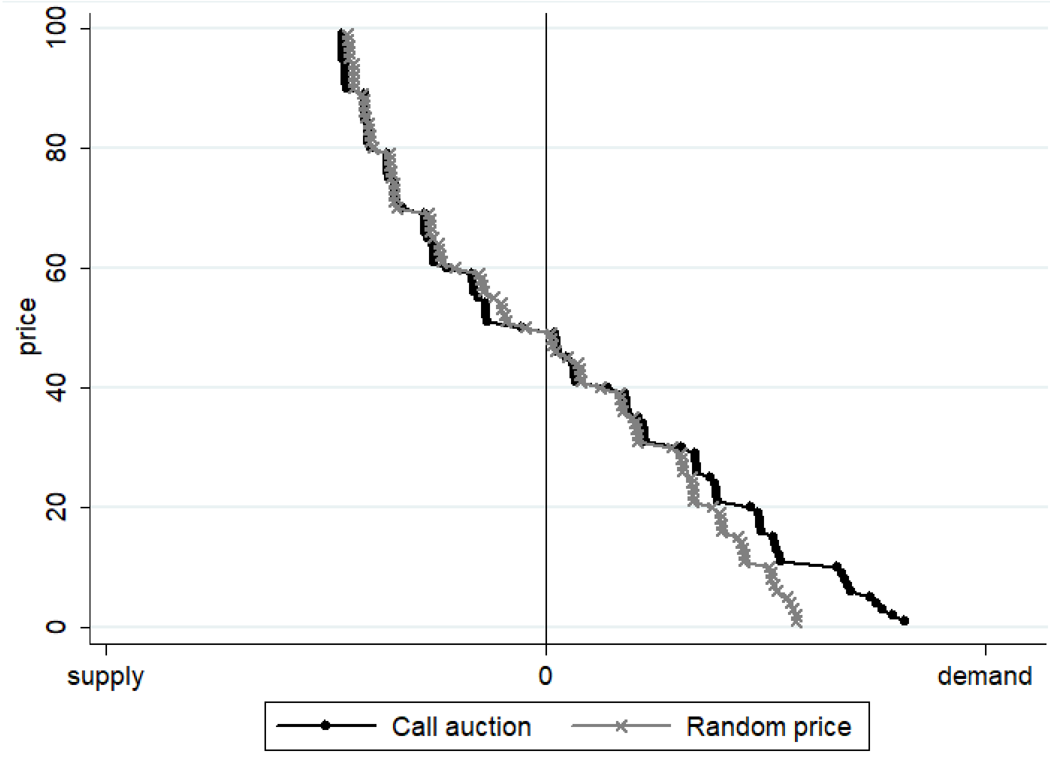

The theoretical model described in Section 2 posits that the degree of risk aversion negatively correlates with the trading volumes at the individual level. In fact, the higher the degree of risk aversion, the steeper the demand, the lower the trading activity at any price. Our experimental results confirm that this is indeed the case. The individual degree of risk aversion θ_{inv} negatively and significantly correlates with the average number of assets exchanged ($\rho = 0.26, p < 0.001$).

Result 1: The degree of risk aversion negatively correlates with the trading volumes at the individual level.

This result is in line with Fellner and Maciejovsky (2007), who also report that measured risk aversion negatively correlates with market activity and is one of the few contributions reporting a significant correlation between elicited risk aversion and market outcomes.

5.3. Elicited risk aversion and market prices

The second testable implication of the theoretical model refers to the role played by risk aversion in the information aggregation process, i.e. in shaping the equilibrium prices. The expected market clearing price is computed under the assumption of homogeneous risk preferences. In order to minimize the variance of risk preferences in every market, we split

Figure 4: Market prices in low and high risk aversion markets.

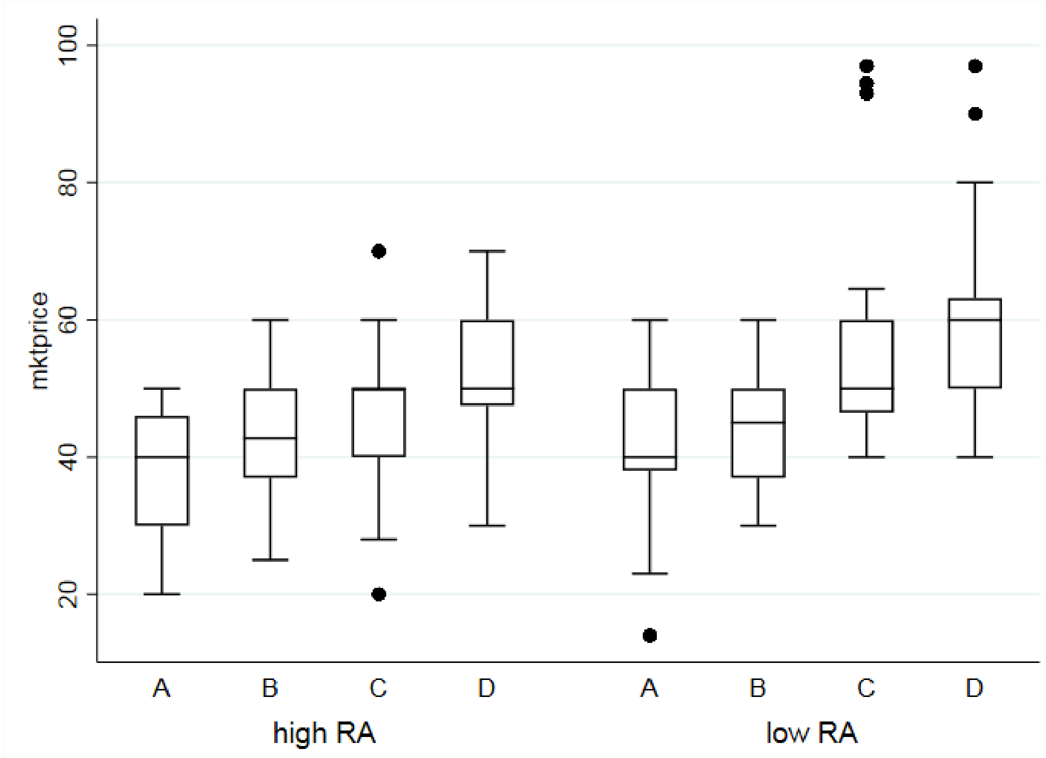

each session between a Low and a High risk-aversion market according to the median choice in the Investment Game in that session.[12]

The average choice in the Investment Game ranges between 40 and 69 in Low risk-aversion markets, and between 119 and 150 in High risk-aversion markets. This manipulation provides an additional test for the effect of the average degree of risk aversion on market prices. As explained in Section 2, the distance of the equilibrium price from the uninformed price (50) should increase with risk aversion. If this is the case, we should observe a steeper pattern of prices in the High vs Low risk aversion markets. Figure 4 shows that this is not the case. Prices contain relevant information, in the sense that they are significantly different between any pair of urns according to a battery of Wilcoxon signed-rank test (using one independent observation per market), but the pattern is similar in High and Low risk aversion markets.

Table 3 compares equilibrium prices in High and Low risk-aversion markets, by testing for differences by urn using a Mann Whitney rank-sum tests with one observation per market. We fail to detect a significant difference for urns A, B, and C. Where a significant

[12] FigureA.13 in Appendix A shows how the heterogeneity of observed risk preferences maps into different predicted equilibrium prices according to the prior information model.

difference emerges, i.e. in urn D, it goes in the opposite direction than predicted under the prior information model because equilibrium prices are higher when risk aversion is lower.

Table 3: Differences between High and Low risk aversion markets

	Urn A		Urn B		Urn C		Urn D	
	U	p-value	U	p-value	U	p-value	U	p-value
Low vs High	-0.907	.364	-.416	.677	-1.566	.109	**-2.307**	**.021**

Notes: the table reports for each urn the Mann Whitney rank-sum test statistic (U), and corresponding p-value, on the difference in equilibrium prices between Low and High risk-aversion markets. A positive statistic means a higher value for High risk-aversion markets. Bold indicates significance at the .05 level. The statistic is computed using one observation per market (20 independent observations, 10 for each of Low and High risk aversion).

The average degree of relative risk aversion in the High risk aversion market is $\theta_{inv} = 1.19$, rather close to log utility. Prices by urn, however, are not close to the average beliefs (29,41,59,71) as the model would predict. In contrast, they are even closer to 50 than in the Low risk aversion markets.

In more detail, Figure 5.3 reports the (absolute) distance of equilibrium prices from the average Bayesian belief against the average CRRA coefficient in each market, estimated using the investment-game decisions. The over-imposed line is a linear fit between the two measures, showing no relation between them, while the expected relationship should be downward sloping and cross the horizontal axes when the CRRA coefficient is equal to 1.

Result 2: The distance between market prices and average beliefs is not minimized around log utility, because equilibrium prices do not react to the average degree of risk aversion in the market.[13]

Given this result, in what follows we analyze the pattern of prices without distinguishing any longer between High and Low risk aversion markets. If the risk aversion of traders does not affect prices, can we conclude that equilibrium prices can be derived under the assumption of risk neutrality? The question is relevant because risk neutrality has long been the standard assumption for the prior information model in experiments on information aggregation in markets (Choo et al., 2017; Plott and Sunder, 1982, 1988). Under risk neutrality, prices should be equal to 45.5 when the urn is A, 50 when the urn is B or C, and 54.5 when the urn is D, while they should be respectively 28.8, 40.9, 59.1 and 71.2 under log-utility. Our average equilibrium prices are in between these two benchmarks, but remarkably closer to the risk neutral one: 40.6 for A, 43.2 for B, 50.6 for C and 56 for D.

The fact that the risk neutral benchmark is a good predictor of equilibrium prices em-

[13]This conclusion is robust to using the average elicited belief, rather than the Bayesian one, as shown by Figure A.12 in Appendix A. In other words, this result cannot be explained by failures of Bayesian updating.

Figure 5: Distance of price from average belief and risk aversion

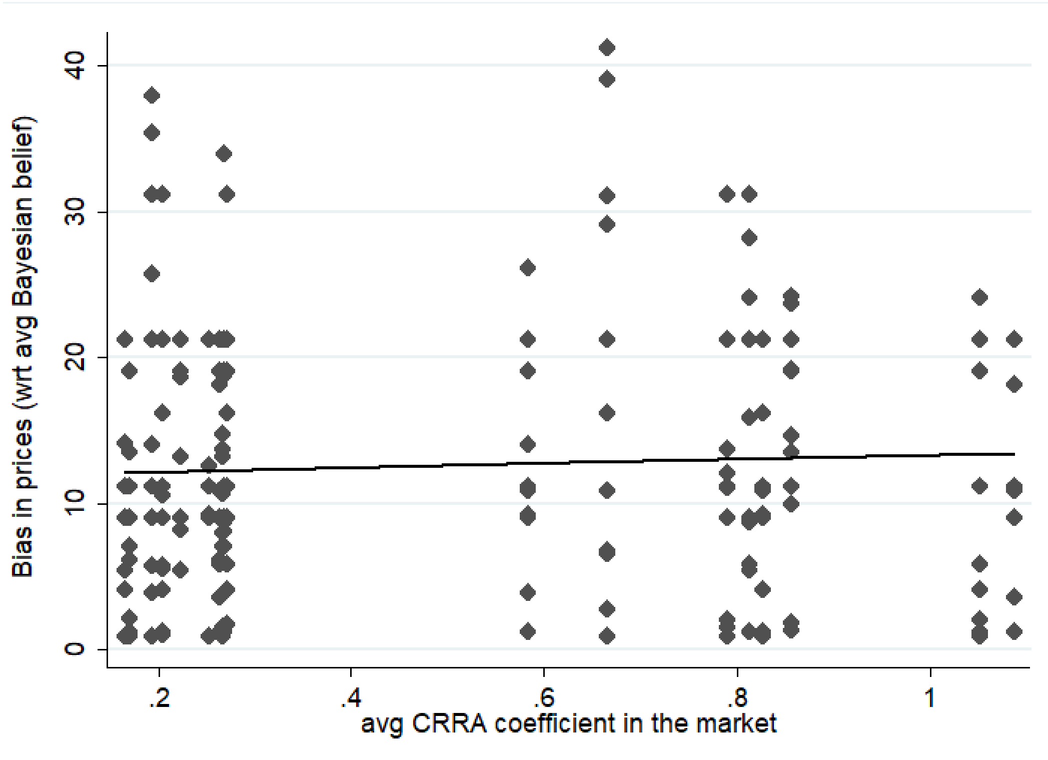

Note: The empty area between average CRRA coefficients of .27 and .58 represents the distance between the maximum average CRRA coefficient among Low risk aversion markets and the minimum coefficient in the High risk aversion markets.

phasizes a contradiction, however. Individuals that show significant levels of risk aversion in the risk elicitation task seem to behave as (approximately) risk neutral at the aggregate level. Indeed, other indicators of behavior in the market induce us to believe that inferring the average degree of risk aversion from the aggregate outcomes could be a misleading exercise. For instance, we know that risk neutral traders should invest all their endowment on bid (ask) orders, whenever the price is below (above) their belief (see Figure 1). Individual demands are clearly at odds with this prediction since only 53.1% of the endowment is committed in trading activity on average. In the next section we delve deeper into this conundrum by estimating individual risk preferences from trading behavior.

5.4. Estimate of individual risk aversion

The experimental setting allows us to estimate the CRRA coefficients from the individual demand schedules. Starting from the net demand function (Equation 2) we input the actual endowment $m = 1000$ and impose the Bayesian beliefs $(0 < b_i < 100)$ consistent with the signal received by each subject in every period.[14] Given that subjects do not extract information from the prices (see Section 5.1) we can exploit the observed individual net demand (q_i) to derive the individual coefficient θ_{mkt} using a Maximum Likelihood estimation.

Figure 3 shows that the observed demand in the CA is roughly linear, which implies that subjects trade systematically less than predicted at extreme values of the prices. A similar phenomenon is reflected by the fact that subjects keep on average 30% of the endowment in their pockets even when the true state of nature is revealed by the signal with certainty, while they should use the whole endowment regardless of their level of risk aversion. Therefore, the estimation is performed using the range of prices between 20 and 80 since including also $p < 20$ and $p > 80$ would overestimate θ_{mkt}.[15]

Not surprisingly, the distribution of individual estimated coefficients displays a pronounced positive skewness given that the lower bound of the coefficient is zero, while extremely high values are sometimes observed. In principle, a subject who never trades would be assigned a degree of risk aversion that tends toward infinity. For this reason, we cap the maximum value of θ_{mkt} to 32.48, i.e. the maximum level attributable within the Investment Game. The average estimated coefficient of risk aversion turns out to be equal to 2.92. Even relying upon the median value, which is not sensitive to outlier decisions, we find that the estimated value (1.86) is of a different order of magnitude than the average elicited one (mean = 0.71, median=0.32). A weird consequence of this finding is that the prices predicted according to the higher estimated θ_{mkt} turn out to be even farther from the observed ones (see Figure 6). Such a larger gap worsens the conundrum, because prices are now predicted using the risk aversion coefficients derived from the choices that generated the observed prices.

The rationale underlying such a counterintuitive finding is illustrated in Figure 7, which shows the average observed net demand by urn, together with the predicted demand that should be expected given the median degree of risk aversion estimated at the individual level. In the theoretical model risk aversion explains the slope of the net demand, i.e. how much subjects are willing to trade as the price moves away from one's beliefs. The higher the risk aversion coefficient, the lower the fraction of the endowment at stake in the market. Figure 7 clearly shows that the estimated values of θ_{mkt} correctly predict the slope of the average net demands. The problem is given by the *intercept* of the predicted market demands, which is lower than the observed one for urn A and B, and higher for urn C and D. Conse-

[14] Using the reported beliefs instead of the Bayesian beliefs delivers similar results, but would impose a sizeable loss of observations.

[15] Note that the model can be estimated only when some uncertainty exists, because the choice should always be $q_i^* = 1000/p$ when $b_i = 100$ and $q_i^* = -1000/(100 - p)$ when $b_i = 0$, regardless of θ_{mkt}. Consequently the number of individual observations ranges from 244 to 732, according to the amount of fully informative signals received.

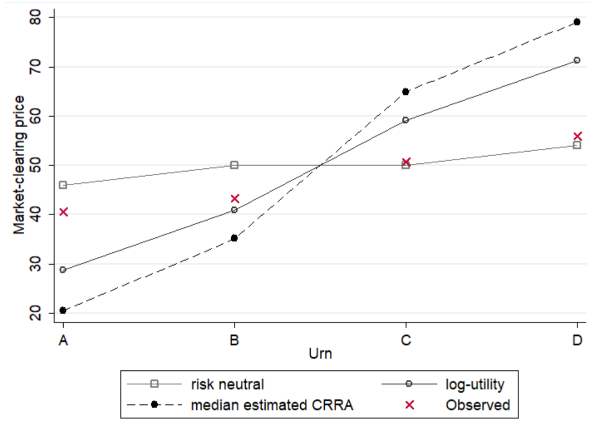

Figure 6: Observed prices are even more inconsistent with estimated risk aversion

quently, the observed aggregate demands always cross the vertical line (and consequently identify an equilibrium price) closer to 50.

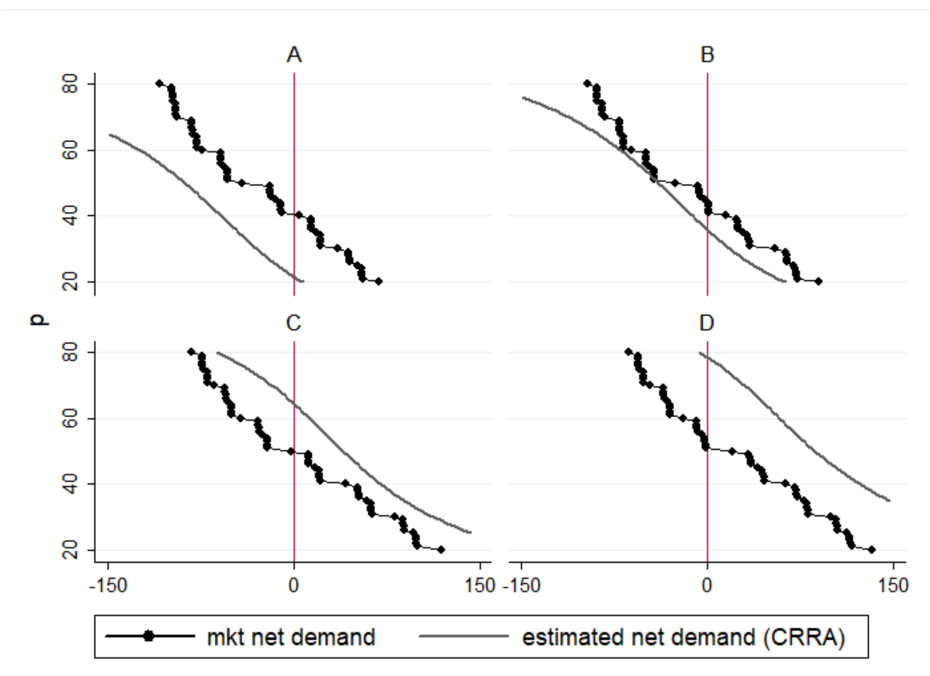

Figure 7: Observed and predicted aggregate net demands

The underlying mechanism is better illustrated analyzing data without aggregating information by signal. Figure 8 displays the average net demand of all subjects when receiving the signal $s = \{53; 54\}$, which implies $b_i = 66.6$, but the argument holds in our data for all the partly or fully informative signals. According to the theoretical model, a trader should switch from buyer to seller around his Bayesian belief. The trader should buy when $p < b_i$

and sell short when $p > b_i$. Figure 8 shows that subjects who know that the asset will pay 100 with two thirds of probability switch instead at a lower price, around 55. In other words, they start selling short below the expected value of the asset.

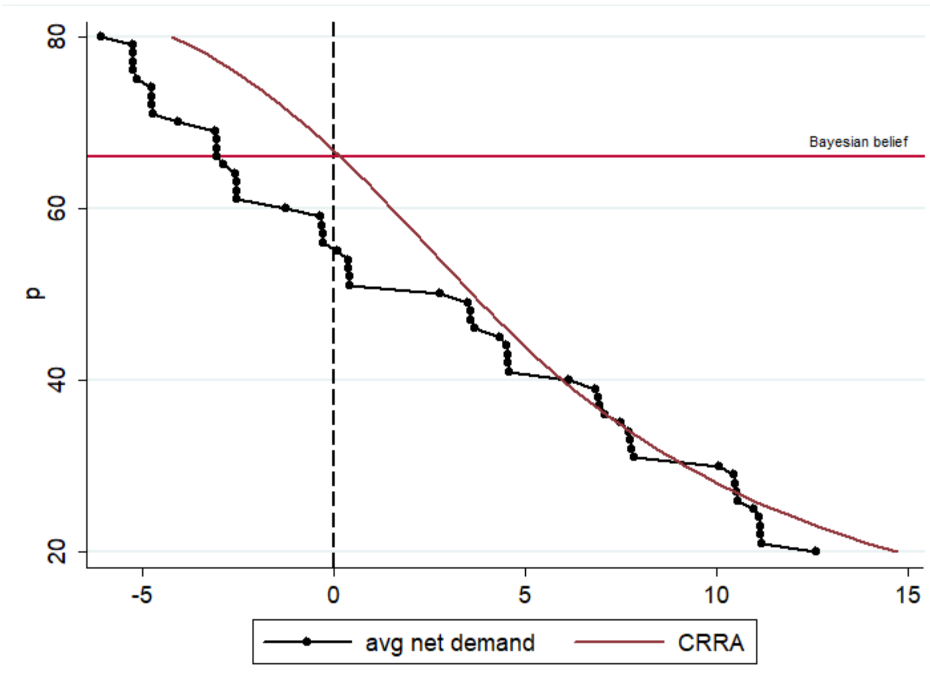

Figure 8: Observed and predicted aggregate net demand for $b_i = 66$

An obvious candidate to rationalize such an aggregate demand is some form of misperception of probability, as if subjects considered the outcome "Majority Blue" as different than it actually was. For instance, conservatism (Peterson and Miller, 1965) posits that variations in subjective probability revision are in the same direction, but of smaller magnitude, than corresponding variations in Bayesian probability change. By making subjects underreact to new information, conservatism would indeed affect the prices in the direction shown by Figure 7. Our experimental protocol allows us to test for the role played by conservatism because we elicit subject's beliefs in an incentivized manner. Table 4 reports the subjects' ex ante beliefs about the state $e = Blue$. Table 4 shows that when receiving fully informative signals subject update their subjective probabilities almost to the maximum extent. In this case mistakes can be made in the direction of conservatism by construction, but the magnitude of the difference is on average remarkably small and beliefs are perfectly correct more than 90% of the times. When receiving partially informative signals subjects instead show evidence of overreaction, as if they consider the true state of nature more likely than what Bayesian update implies. Overall, choices are not driven by conservatism, at least in the way in which this behavioral trait has been defined.

Signal	Bayesian	Reported
42-45	0	3.8
46-47	33.3	27.3
48-52	50	52.4
53-54	66.6	75.6
55-58	100	97.1

Notes: The subjective probability of $e = Blue$ is computed summing up the probability assigned to urn C and D.

Table 4: Ex ante Beliefs about the state $e = Blue$

The fact that switching from buyer to seller occurs at prices closer to 50 reveals that the information about the true state of nature, although fully internalized, is transferred to the choices only in part. In other words, the divergence between predicted and observed prices can be accounted for by subjects behaving *as if* they are less informed. Indeed, a trader who use his information only in part is behaviorally indistinguishable from another trader who partially updates his subjective probability and then acts fully exploiting such conservative beliefs. On the other hand, it is worth noting that while in the theoretical model beliefs map straightforwardly into the intercept of the net demand, the experimental task is cognitive demanding and such a direct link cannot be assumed in the mind of the subjects. Inserting buy and sell orders is not like stating a subjective probability and subjects may easily miss the theoretical connection between the two things. This is another reason why we deem semantically more appropriate to refer to this phenomenon as a reluctance to act according to the information possessed rather than to conservatism.

Following a formalization similar to that of Epstein (2006) we assume that subjects behave according to a belief $\hat{b}_i$, which is a convex combination of the information actually received b_i and the uninformed state ($b_i = 50$):

$$\hat{b}_i = (1 - \delta_i)b_i + \delta_i(50). \tag{5}$$

The new specification includes a parameter δ_i meant to capture the amount of information not incorporated into the choices. We then re-estimate the individual demand (Equation 2) substituting b_i with $\hat{b}_i$. Note that δ_i is estimated only when subjects receive an informative signal, while the parameter is not affected by the choices when $s \in [48, 52]$. The effect of δ is that of shifting the net demand, with the intercept moving toward 50 as $\delta_i \to 1$.

The two-parameter (θ_{mkt}, δ) version of the model delivers several interesting results. First, the estimate of δ displays that a lot of information is not used. This parameter turns out to be significantly larger than zero for 141 subjects out of 219, with an average value equal to 0.64.[16] Given the complexity of a market experiment, one could also argue that the waste

[16]The number of orders of one subject was insufficient to estimate his parameters.

of information is a short run phenomenon, which should disappear as long as the subjects gain sufficient experience. Modelling a non-linear learning process in the estimation of δ at the aggregate level, we indeed find that the waste of information decreases over time, but the magnitude of learning is small and shrinking. The model predicts the convergence to a long run value of $\delta = .50$.[17]

Second, the distribution of estimated θ_{mkt} changes slightly when δ in included in the model. The median estimated value of θ_{mkt} decreases from 1.86 to 1.79. Despite the small magnitude of the change the two distributions differ significantly according to a Wilcoxon signed-rank test ($z = 2.469$, $p = 0.014$) suggesting a degree of substitutability between θ_{mkt} and δ that will be discussed in more detail in Section 5.5.

Third and foremost, the (θ_{mkt}, δ) specification allows us to properly reconstruct the aggregate behavior in the market, and therefore to reconcile predicted and observed equilibrium price, as shown by Figure 9. The gap between observed and predicted demands disappears almost completely when the model accounts for a partial use of the available information at the individual level.

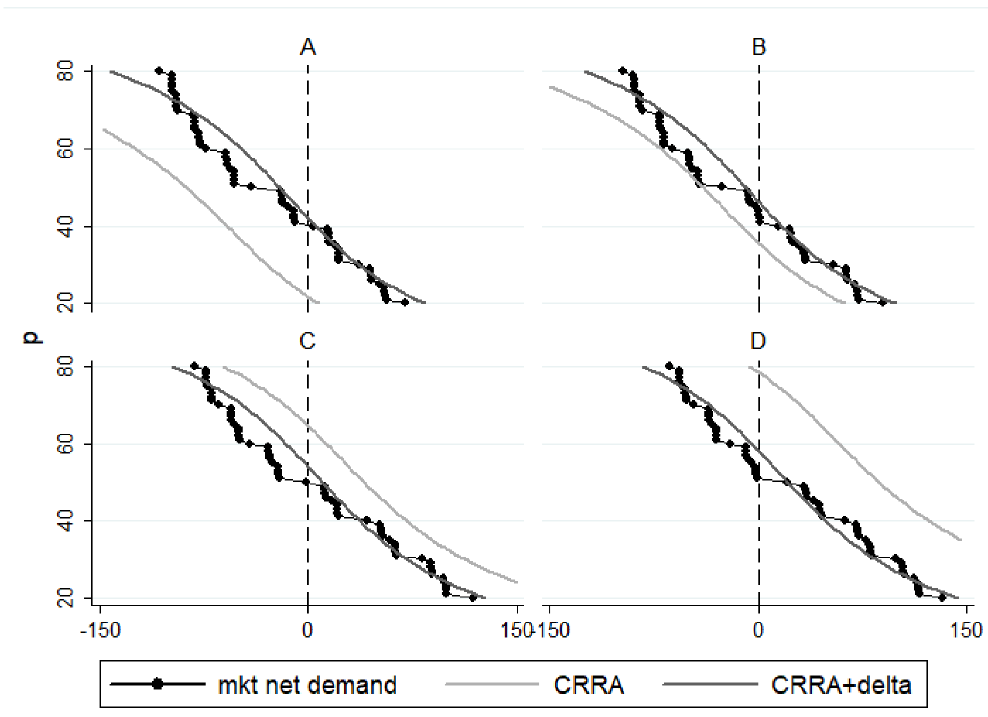

Figure 9: Observed and predicted (with and without δ) aggregate demands

An important corollary of this exercise is to emphasize why trying to infer the average

[17]Results are available upon request.

degree of risk aversion in the market from the equilibrium prices can be a misleading exercise. Even when the amount and the distribution of information in the market is known, we cannot take for granted that it will be fully exploited. In our experiment, the observed pattern of prices would aim at risk neutrality, while the exercise just performed show that the pattern of prices reflects instead the aggregation of a lower amount of information than what initially distributed.

An extremely interesting feature of $\delta > 0$ is that it can behaviorally be interpreted in terms of risk aversion, since it induces a lower exposure in the market. Figure 10 describes the estimated aggregate demand with and without δ when the underlying urn is A, but the same argument applies to the other urns as well. A positive value of δ implies a net demand that features no trade closer to 50 than around the Bayesian belief of 28.9. The upward shift of the demand causes more buy orders than what should be optimally observed (dark shaded area), but an even larger reduction of short selling (light shaded area). The net effect of $\delta > 0$ is clearly a less intense trading activity.

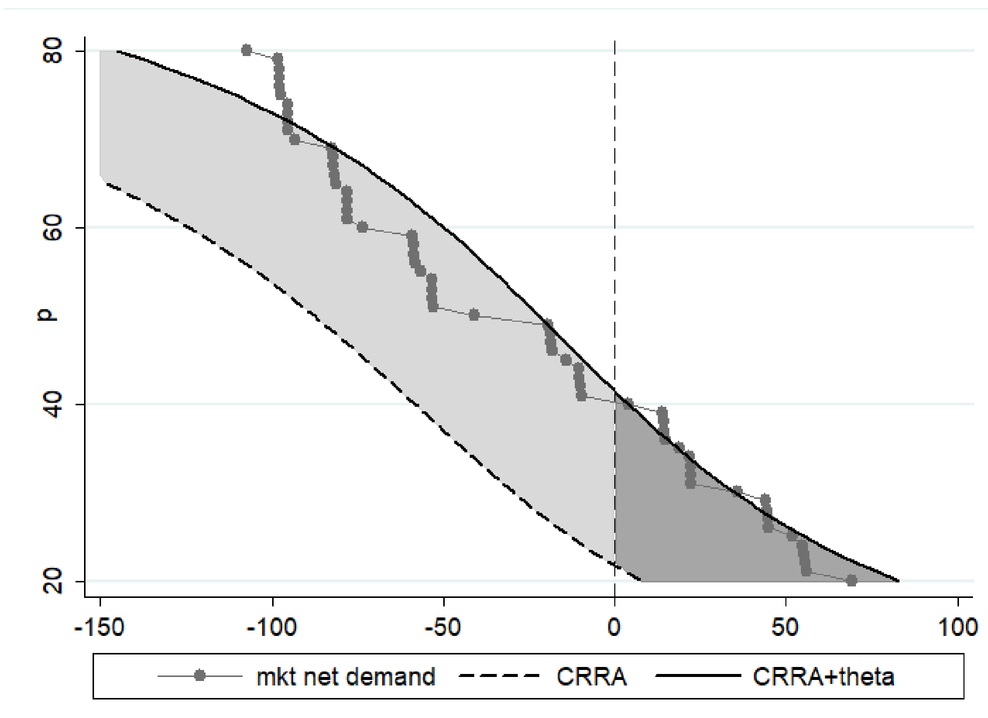

Figure 10: Delta and market exposure

Although buying and short selling are equally risky, the waste of information can represent an additional facet of a "reluctance to act" besides that represented by the slope of the net demand driven by θ_{mkt}. Data at the individual level confirm such an interpretation: the standard deviation of potential earnings in all the 12 periods negatively and significantly correlates with δ ($\rho = -0.136$, $p = 0.045$). On the other hand, the waste of information implies a suboptimal behavior and in fact the average potential earnings negatively and strongly correlate with δ ($\rho = -0.50$, $p <= 0.001$).[18] In other words, $\delta > 0$ constitutes an additional mechanism through which subjects can reduce the variance of the outcomes at a (inefficiently high) cost, and this is the reason why it can be interpreted in terms of risk aversion.

5.5. Consistency of risk aversion measures

Moving to the individual measures of risk aversion we now have two values for every subject, the first elicited with the Investment Game (θ_{inv}), and the second estimated using the behavior in the market (θ_{inv}). The first thing to notice is that the elicited and estimated measures are of a different order of magnitude. The average (median) CRRA coefficients are $\theta_{inv} = 0.71(0.32)$ and $\theta_{mkt} = 2.92(1.78)$, respectively. Such a pronounced difference may signal that subjects perceive the market as more risky institution. Possibly, the virtually infinite set of outcomes without objective probabilities implied by the call market induces a more prudent behavior than the binary lottery with equally likely outcomes in the Investment Game.

Apart from the different levels, a natural question is to check whether the two measures are consistent with each other. That individuals' risk preferences may not be stable across elicitation methods, indeed, would come at no surprise given the plethora of contributions already showing that this is the case (see Deck et al., 2013; Friedman et al., 2014; Isaac and James, 2000, among others). In what follows we show that our results do not support such a negative conclusion. The linear correlation of the CRRA coefficients is not significant ($\rho = 0.08$, $p = 0.216$). This result is not surprising, however, because it is driven by the presence of many outliers since the distributions of the two parameters are characterized by a pronounced positive. A significant correlation emerges, in fact, when the effect of the outliers is removed using the ranks (Spearman's $\rho = 0.16$, $p = 0.015$). Another exercise can be performed comparing θ_{mkt} with the rough choice in the Investment Game rather than with θ_{inv}. As already explained in Section 5.4, choices map onto θ_{inv} in a highly non-linear manner, with investments of at least half of the endowment that have a disproportionately lower effect in shaping θ_{inv}. In this case the linear correlation turns out to be significant ($\rho = -0.18$, $p = 0.006$).[19]

Given the interpretation of δ in terms risk aversion, it is interesting to check whether also this aspect is captured by the representation of decisions under risk hold by the sub-

[18] The significant correlations hold even when controlling for the individual θ_{mkt}.

[19] The correlation in this case has a negative sign because the higher the amount invested, the lower the degree of risk aversion.

jects. In other words, subjects may express their risk aversion not having in mind simply the curvature of the utility function (as economists do), but instead as a broader construct encompassing a more general "reluctance to act." The idea is that subjects' representation of risk aversion may have two distinct dimensions connected to the parameter θ and δ. The first is the intensity of trading given beliefs, while the second is the inhibition to act fully exploiting the information available. Indeed, the choice in the Investment Game correlates significantly with δ ($\rho = -0.18$, $p = 0.006$). The higher the amount invested in the risk elicitation task, the lower δ, i.e. the less reluctant is the subject to act consistently with the information received. Table 5 shows that elicited risk aversion significantly correlates with θ_{mkt} and δ both separately (Column 1 and 2) and at the same time (Column 3).

Table 5: Consistency of measures of risk aversion

	(1)	(2)	(3)	(4)
	Dependent var: Choice in the Investment Game			
θ_{mkt}	-1.655***		-1.766***	-1.783***
	(0.501)		(0.485)	(0.497)
δ		-20.76***	-21.57***	-20.00***
		(5.173)	(4.718)	(5.330)
Errors				-0.160
				(0.783)
Low financial literacy				-7.270
				(7.096)
Constant	98.28***	106.7***	112.4***	115.5***
	(9.926)	(10.34)	(11.10)	(12.29)
N	219	219	219	219
R^2	0.019	0.034	0.056	0.061

Robust standard errors. Data clustered at the market level. Significance $^{***} = 0.01$.

Since holding $\delta > 0$ is costly, one could reasonably argue that such a reluctance to act may characterize the subjects with a bad understanding of the market mechanism. The fourth column of Table 5 includes the errors made by the subjects in the quizzes and a dummy capturing their (self-reported) low degree of financial literacy. The coefficients are robust to the inclusion of the additional controls, showing that the reluctance to use the information is not driven by a more limited knowledge of how to behave.

This result is intriguing in our opinion, as it helps explaining why elicited risk attitudes usually have a very limited predictive power. While economists restrict the risk aversion concept to the diminishing marginal utility of money, subjects likely hold a broader representation of this construct. Such an interpretation is indirectly suggested by the whole literature displaying instability and inconsistency in the measurement of risk preferences. However, to our knowledge Table 5 constitutes the first piece of evidence identifying in a solid manner an additional determinant within their representation beyond loss aversion.

The broader perspective including the reluctance to exploit available information also allows us to reconcile the individual characteristics with the market outcomes. We have seen before (Figure 4) that market prices do not differ between High and Low risk aversion markets, while prices were expected to react more to the urn in the first case. The positive correlation between θ_{inv} and δ at the individual level shown in Table 5 reveals the effect that counterbalances that of the higher risk aversion. Subjects in the High risk aversion markets should also be more reluctant to exploit the information available, pushing market prices closer to 50 . The data confirm this intuition: δ is significantly higher (0.70 Vs. 0.58, $p = 0.061$, respectively) in the High risk aversion markets.

At the aggregate level the estimated coefficients of risk aversion are also significantly higher in the High risk aversion markets (median 1.95 Vs. 1.72, $p = 0.078$). However, θ_{mkt} and δ display a remarkable degree of substitutability at the individual level (Spearman's $\rho = -0.27$, $p < 0.001$). The interpretation is that a more pronounced reluctance to act is what characterizes the subject in the High risk aversion market, as correctly captured by the choices in the Investment Game, but the reluctance to act is mainly expressed either through a steeper net demand or exploiting less the information available.

6. Conclusion

In this paper we analyse the effect of risk aversion on information aggregation in call markets. After checking that subjects do not extract information from the conjectured equilibrium prices, we find that the informational content of realized equilibrium prices – in terms of representing traders' beliefs – does not reflect the average degree of risk aversion in the market. Contrarily to the theoretical prediction, markets characterized by significantly different risk aversion deliver indistinguishable patterns of prices. In a parallel manner, it is impossible to infer traders risk aversion from aggregate market outcomes. While prices are close to predictions under risk-neutrality, risk attitudes as both independently elicited and estimated from individual net demands point toward a significant degree of risk aversion.

We investigate the mechanism through which the interaction among risk-averse traders may lead to risk-neutral prices. By estimating a two-parameter $(\theta; \delta)$ model fitting the net individual demands under CRRA we find that two different mechanisms shapes traders' behavior. The first (θ) is the slope of the net demand, i.e. how much the net demand reacts to changes in the price of the asset, which can be attributed to the classic risk aversion coefficient. The second (δ) is instead the intercept of the net demand, which represents the inclination of subjects to exploit the information possessed. These two facets of cautious behavior have opposite affects on market prices, which explains why markets composed of players with a different degree risk-aversion end up showing similar market-clearing prices. Moreover, the parameter δ allows us to emphasize that the information aggregated by the market is only part of the information possessed. By aggregating a lower amount of information, prices turns out to be observationally similar to those that would have been observed under risk neutrality and full aggregation of information.

A very interesting result in our opinion is that both components significantly correlate with a measure of risk aversion obtained with the Investment Game. Elicited risk aversion seems to capture a 'reluctance to act' that is more general than the classic curvature of the utility function. This finding delivers an original and important message to the literature on the measurement of risk preferences. When models (and risk elicitation methods) intend risk-aversion only as the curvature of the utility function, they may fail to capture important features of what subjects represent with their choices as a risk-averse behavior. Overall, our results do not add to the long list of contributions claiming the empirical failure of the measurement of risk preferences when not of expected utility theory as a whole. Our evidence is more optimistic than the consensus in this branch of the literature about the the possibility to capture stable features of choice under uncertainty with simple elicitation tasks. Risk preferences are a latent construct that even in the best case scenario can be observed only together with a huge measurement error. At least when such a measurement error is sufficiently reduced (e.g. by restricting the domain of choices to the financial sphere, by avoiding confounds such as loss aversion, by choosing an elicitation method that implies a kind of choices comparable to that of the main task) consistent choices can be observed even in a complex environment like a call market, particularly when risk attitudes are not restricted to the curvature of the utility function in evaluating the variance of the outcomes.

References

Alti, A., Tetlock, P. C., 2014. Biased beliefs, asset prices, and investment: A structural approach. The Journal of Finance 69 (1), 325–361.

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-new york then washington- 320 (5878), 877.

Becker, G., DeGroot, M., Marschak, J., 1964. Measuring utility by a single-response sequential method. Behavioral Science 9, 226–236.

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

Charness, G., Gneezy, U., 2012. Strong Evidence for Gender Differences in Risk Taking. Journal of Economic Behavior & Organization 83 (1), 50–58.

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

Cipriani, M., Guarino, A., 2014. Estimating a structural model of herd behavior in financial markets. American Economic Review 104 (1), 224–51.

Crosetto, P., Filippin, A., August 2013. The 'bomb' risk elicitation task. Journal of Risk and Uncertainty 47 (1), 31–65.

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

Deck, C., Lee, J., Reyes, J. A., Rosen, C. C., 2013. A failed attempt to explain within subject variation in risk taking behavior using domain specific risk attitudes. Journal of Economic Behavior & Organization 87 (0), 1 – 24.

Easley, D., Kiefer, N. M., O'Hara, M., 1997. One day in the life of a very common stock. The Review of Financial Studies 10 (3), 805–835.
URL `http://dx.doi.org/10.1093/rfs/10.3.805`

Epstein, L. G., 2006. An axiomatic model of non-bayesian updating. The Review of Economic Studies 73 (2), 413–436.

Fellner, G., Maciejovsky, B., 2007. Risk attitude and market behavior: Evidence from experimental asset markets. Journal of Economic Psychology 28 (3), 338 – 350.

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

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

Friedman, D., Isaac, R. M., James, D., Sunder, S., 2014. Risky curves: On the empirical failure of expected utility. Routledge.

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.

He, X.-Z., Treich, N., Jan. 2013. Heterogeneous Beliefs and Prediction Market Accuracy. LERNA Working Papers 13.05.392, LERNA, University of Toulouse.

Holt, C., Laury, S., 2002. Risk aversion and incentive effects. American Economic Review 92 (5), 1644–1655.

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

Isaac, R., James, D., 2000. Just who are you calling risk averse? Journal of Risk and Uncertainty 20 (2), 177–187.

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

Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. Journal of Economic Theory 26 (1), 17–27.

Ngangoue, K., Weizsacker, G., 2015. Learning from unrealized versus realized prices (1487).

Peterson, C. R., Miller, A. J., 1965. Sensitivity of subjective probability revision. Journal of Experimental Psychology 70 (1), 117.

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.

Plott, C. R., Sunder, S., 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. Econometrica: Journal of the Econometric Society, 1085–1118.

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. Further results

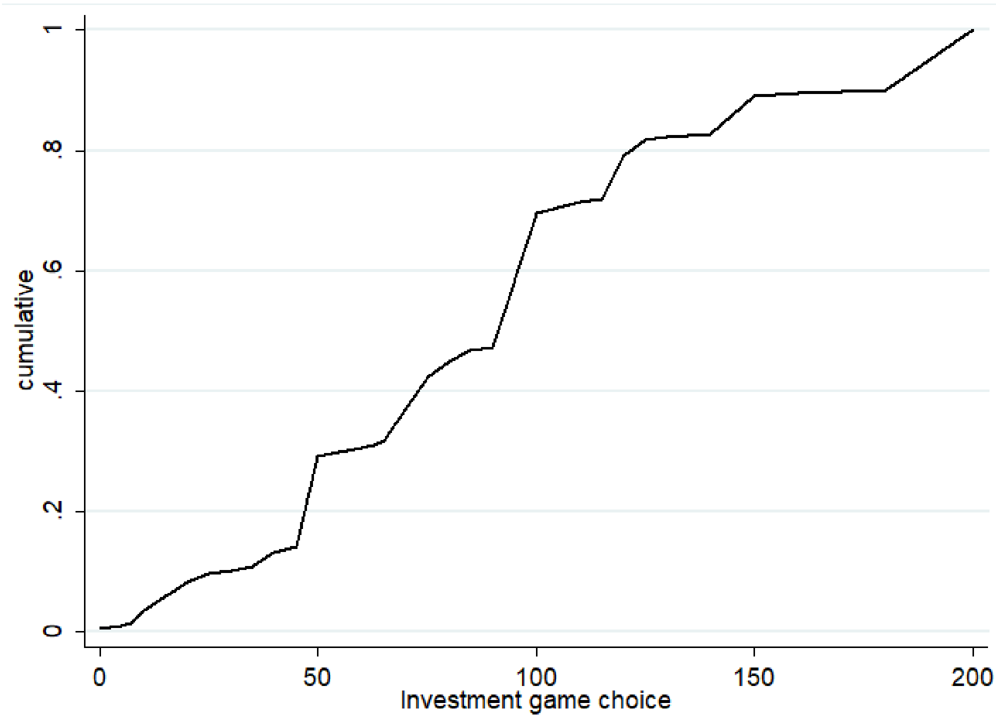

Figure A.11: Cumulative distribution of choices in the risk elicitation task

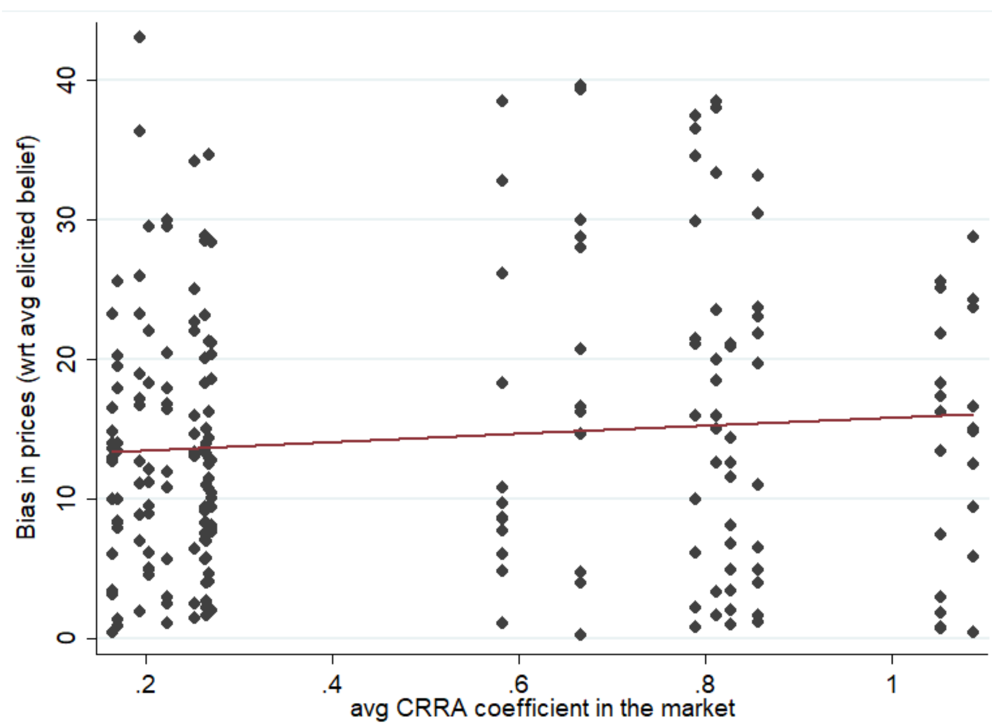

Figure A.12: Distance of price from average (elicited) belief and risk aversion

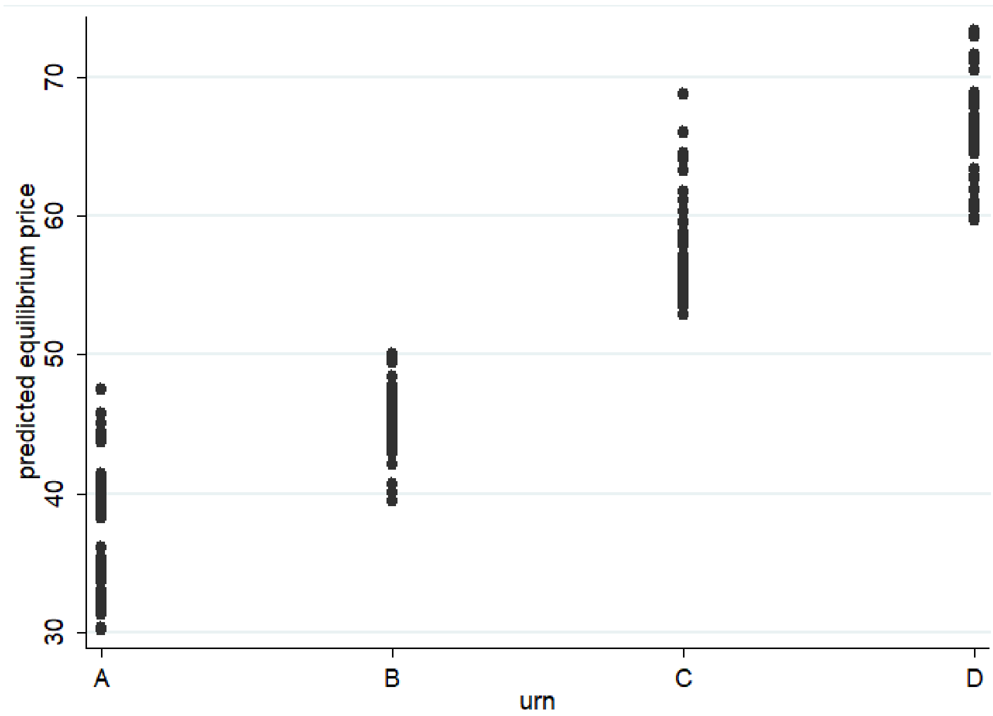

Figure A.13: Predicted equilibrium prices, given the joint distribution of signals and CRRA