

Filippin, Antonio; Mantovani, Marco

Article

Risk aversion and information aggregation in binary-asset markets

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Filippin, Antonio; Mantovani, Marco (2023) : Risk aversion and information aggregation in binary-asset markets, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 14, Iss. 2, pp. 753-798, <https://doi.org/10.3982/QE1981>

This Version is available at:

<https://hdl.handle.net/10419/296330>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Risk aversion and information aggregation in binary-asset markets

ANTONIO FILIPPIN

Department of Economics Management and Quantitative Methods, University of Milan and Institute for the Study of Labor (IZA)

MARCO MANTOVANI

Department of Economics, Statistics and Management, University of Milan Bicocca

We investigate how risk aversion (RA) shapes the informative content of prices in an experimental asset market, where traders are sorted according to their RA. RA should induce steeper individual demands and, under its most common parametrizations, drive equilibrium prices closer to revealing the state. Results support the prediction on individual demands, but not the prediction on prices, which do not vary with RA and are close to the risk-neutral benchmark. This purported conflict is due to traders, particularly the more risk-averse ones, conveying into prices only part of their information.

KEYWORDS. Risk preferences, laboratory experiment, asset markets, information aggregation, Walrasian equilibrium, operational conservatism.

JEL CLASSIFICATION. C92, D81, G14, G41.

1. INTRODUCTION

Markets play a central role in aggregating dispersed information, and their performance on this dimension is the main yardstick to evaluate their efficiency. As long as prices are expected to reflect traders' information, extracting information from prices becomes a natural exercise, particularly in financial markets (e.g., [Easley, Kiefer, and O'Hara \(1997\)](#), [Alti and Tetlock \(2014\)](#), [Cipriani and Guarino \(2014\)](#)). Some specific asset markets, called prediction markets, exist with the unique purpose of inferring beliefs from prices (e.g., [Wolfers and Zitzewitz \(2004\)](#), [Arrow et al. \(2008\)](#)). Prediction markets are increasingly

Antonio Filippin: antonio.filippin@unimi.it

Marco Mantovani: marco.mantovani@unimib.it

We would like to thank three anonymous referees as well as Peter Bossaerts, Dan Friedman, Anna Gilbert Rivas, Stephanie Heger, Peter Katuščák, Thomas Kittsteiner, and participants to seminars at the University of Florence, RWTH Aachen University, University of Milan, and the University of Innsbruck as well as participants to the CEREC Workshop 2018 (Bruxelles), the Experimental Finance Conference 2018 and 2019 (Heidelberg, Copenhagen), the WZB workshop "The Neoclassical Repair Shop and its Alternatives" (Berlin, 2019), and the MBEES workshop 2019 (Maastricht) for useful suggestions. All remaining errors are ours. We gratefully acknowledge financial support from IFREE Small Grant Program. Marco Mantovani gratefully acknowledges financial support from the Italian Ministry of University and Research, grant PRIN 2017K8ANN4. Data, materials and replication files: <https://osf.io/d6ey5/>. A replication file is posted ([Filippin and Mantovani \(2023\)](#)).

used by corporations to support the innovation management (Horn and Ivens (2015)) as well as a forecasting device for political, economic, and sport events (see, for instance, Luckner, Schröder, and Slamka (2008), and references therein).

Except in the limit case in which no uncertainty survives, as in a fully revealing Rational Expectation Equilibrium (REE) without aggregate risk, the link between information and prices is crucially mediated by the risk preferences of the traders. Notwithstanding their crucial role theoretically, and despite trading being an archetypical instance of decision under uncertainty, the empirical evidence on how risk preferences shape information aggregation is scant.

The goal of this paper is to investigate how risk preferences shape the information aggregation process. Manski (2006) and Gjerstad (2005) show how risk aversion affects the relation between the market price and the traders' beliefs in a setting with an exogenous belief distribution. Assuming a Prior Information Equilibrium (PIE), we can extend their results to derive predictions in a context where beliefs reflect information on the state of the world.

In a PIE, naive traders simply bring their private information to the market without making any inference about their opponents' information. In this framework, the more risk averse are the traders, the closer to full revelation of the state is the market-clearing price, both under constant absolute (CARA) and constant relative risk aversion (CRRA). To grasp the underlying intuition, note that a risk-neutral trader bets the whole endowment on short (long) positions whenever the expected value of the asset is below (above) the price. Therefore, his demand does not respond to the distance between the price and the expected value of the asset given his beliefs or, in other words, to how precise is the information he holds. A risk-averse trader, instead, bets a larger fraction of his endowment when he holds more precise information about the value the asset. Risk-averse demand schedules, and consequently market prices, therefore, are more responsive to information than risk-neutral ones.

The institution that better fits the PIE is the single call auction in which subjects place limit bid and ask orders in a closed book. The aggregation of individual demand schedules identifies a market-clearing price, with all the exchanges occurring at that price. In our 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 meant to induce heterogeneous posterior beliefs. Since failures in Bayesian updating may influence prices, we elicit subjects' beliefs in an incentive compatible manner. Given the state, we induce common preferences over the asset and we can isolate different information as the only driver of trade. Note that even though trades are zero-sum, the no-trade theorem (Milgrom and Stokey (1982)) does not hold under the PIE. In this way, our design allows us to identify risk aversion as the unique mediating variable in the theoretical model.

Prior to entering the market, we elicit traders' risk preferences using the Investment Game (Gneezy and Potters (1997)). We use the median choice in this task to divide the traders in each session in two groups.¹ That is, we exogenously induce markets with substantially different risk-aversion levels. Across these markets, the level of risk aversion

¹Crockett, Friedman, and Oprea (2021) sort subjects based on risk preferences showing that it helps to predict the outcome of general equilibrium exchange economies.

should have an unambiguous effect on market prices (at least under CARA and CRRA): prices should be closer to full revelation of the state in the more risk-averse markets. Moreover, we manipulate exogenously the amount of information distributed across periods within each market.

To our knowledge, this paper provides the first thorough investigation of the relation between the risk preferences of the traders and the information aggregation properties of an experimental asset market. Some contributions detect risk aversion in experimental markets with aggregate risk, by showing that risky assets are traded at prices below their expected value (Bossaerts and Zame (2008), Biais, Mariotti, Moinas, and Pouget (2017)).² Other studies rule out the role of risk aversion investigating information aggregation in a REE without aggregate risk. Under PIE, however, risk preferences matter even without aggregate risk, but the mechanism through which they shape the aggregation of information has not been directly tested.

The existing evidence on call markets supports the assumption that when deciding how much to buy or sell at any price, traders do not internalize the information brought by the other traders if that price happened to be the market-clearing one. Ngangoué and Weizsäcker (2021) find that subjects are unable to conjecture the informational content of prices through hypothetical reasoning when the mechanism requires to submit orders conditional on the price (as in a call auction). Biais et al. (2017) compare a call auction and a random price mechanism to test for the assumption of competitive behavior in a complete market. They fail to detect a difference between the two mechanisms. In a large meta-analysis of binary-asset markets, Page and Siemroth (2021) find little aggregation of private information even in double auctions. However, in our set-up the effect of strategic behavior would be confounded with that of risk aversion. Therefore, we specifically test the goodness of the PIE assumption by comparing the demands submitted in the call market with those observed under a random-price mechanism, in which prices cannot depend on traders' information by construction. We find that demands are indistinguishable between the two institutions, supporting the validity of the prior information model to derive our testable implications about the role of risk aversion.

We find that choices in the risk elicitation task predict behavior in the market: more risk-averse individuals demand and supply significantly lower quantities. However, risk attitudes are not reflected by equilibrium prices across the board. First, we do not observe significant differences between low and high risk-aversion markets. Second, prices are weakly informative, in the sense that they feebly respond to the overall amount of information given to traders. The observed pattern of prices resembles the prediction under risk neutrality, while choices in the Investment Game display a significant degree of risk aversion.

²Other studies emphasize significant correlations of risk aversion with specific aspects of the markets. For instance, Fellner and Maciejovsky (2007) find that risk aversion, as measured through binary lottery choices, predicts market activity (bids, asks and trades) in double auction markets. Pennings and Smidts (2000) show that elicited risk aversion is a better predictor than a psychometric measure also for market decisions in the field. Ang and Schwarz (1985) compare one low and one high risk-aversion market. They find higher risk premiums in the latter, combined with lower price volatility and lower efficiency.

To uncover the mechanism leading risk-averse traders to exhibit seemingly risk-neutral prices, we estimate individual risk-aversion parameters from the individual demands schedules. Indeed, estimated risk-aversion coefficients are larger than those elicited with the Investment Game and, therefore, even farther from risk neutrality. While these parameters correctly capture the slope of net demands, observed demands schedules are shown to be shifted in the direction of a less informed behavior than the predicted ones. Individuals act as if they had less information than they actually possess, according to both the Bayesian and their elicited beliefs. Consequently, we reestimate the individual risk-aversion coefficients incorporating in the net demands a second parameter capturing this *operational conservatism*.

The two parameter model allows us to rationalize the experimental results at the market level. On the one hand, *operational conservatism* drives market-clearing prices in the opposite direction as compared to classic risk aversion, explaining why we do not find differences between low and high risk-aversion markets. On the other hand, *operational conservatism* induces less informative prices, rationalizing why prices react less than predicted to the amount of information available in the market. In other words, the observed pattern of prices reflects a lower amount of information than that actually distributed and received, rather than risk neutrality.³

We show that *operational conservatism* reduces the exposure in the market and constitutes another way to trade lower expected returns with a lower variance of earnings. Thus, while individuals have a more efficient way of doing so (by submitting a steeper demand schedule), *operational conservatism* may represent an additional (and sub-optimal) way of expressing one's attitudes toward risky environments. We then check whether the reluctance to act exploiting the information possessed, as well as the estimated coefficients of risk aversion, correlate with the choices in the risk elicitation task. We find indeed a positive and significant correlation in both cases, suggesting that individuals express a cautious behavior in different ways. The correlation among the three measures could be driven by common factors other than risk aversion, such as lack of understanding or confidence in one's beliefs. While we cannot rule out this possibility, we show that this correlation is not driven by the inability to process information as captured by errors in the quizzes and self-reported financial literacy.

The paper proceeds as follows. Section 2 discusses the theoretical role of risk aversion and derives the main testable implications. Sections 3 and 4 present the experimental design and the procedures, respectively. Results are reported in Section 5. Section 6 concludes.

2. THE THEORETICAL ROLE OF RISK AVERSION

We study the role of risk aversion using a common-value asset market with heterogeneous beliefs and without aggregate risk, that is, the theoretical framework commonly used to model prediction markets. There are N traders facing uncertainty regarding two ex ante equally likely states, $e \in \{Red, Blue\}$. An Arrow–Debreu security is traded on a

³As a corollary, our results speak against inferring risk preferences directly from market prices (e.g., Cox, Roberson, and Smith (1982), Bliss and Panigirtzoglou (2004)).

market. The security pays 100 to its owner if $e = \textit{Blue}$, and pays 0 if $e = \textit{Red}$. The price of the security will then be $p \in [0; 100]$.

Each trader forms a belief $b_i \in [0, 100]$ exploiting his private information, where b_i represents his subjective probability that $e = \textit{Blue}$, in percentage points. That is, for each trader, b_i coincides with the expected value of the asset. In this section, we take b_i as given.⁴

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. For the sake of simplicity, we refer to m as “monetary endowment.” Since there are no endowments of the security, sales occur through short selling. Short positions are covered at the closure of the market at the actual value of the security, given the realized state. That is, sellers buy back the asset at a price equal to 100 if $e = \textit{Blue}$ and to 0 if $e = \textit{Red}$.⁵

In a Single Call Auction (henceforth: CA) traders submit a demand schedule $q_i(p)$, for $p \in [0, 100]$, where a negative demand at a given price indicates a short position. 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 $p^* : \sum_i q_i(p^*) = 0$.⁶ We assume that traders behave as expected utility maximizers and act as price-takers on the market.

2.1 *Prior information model*

We derive our hypotheses under the so-called *prior information equilibrium*. 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 to the market, but do not extract information from it before the market-clearing price is revealed. Therefore, prices simply aggregate the information the traders have *prior* to entering the market and their behavior is nonstrategic.⁷

We derive our hypotheses using the CRRA specification as a working example. In Appendix B, we show that all of our comparative statics on risk aversion (Hypotheses 1,

⁴A perhaps more standard scale would have the asset paying 0 or 1, $b_i \in [0, 1]$ and $p \in [0, 1]$. The rescaling, where probabilities are written in percentage points, keeps this section consistent with the experimental design. In Section 3, we discuss how beliefs are induced through informative signals and how they are elicited in the experiment.

⁵This 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 100 when $e = \textit{Red}$. Considering a single security simplifies the experimental task for the subjects.

⁶We assume for the moment that such a price exists and is unique. In footnote 15, we explain how we deal with imperfect market clearing and with multiple market-clearing prices in the experiment.

⁷The effects of relaxing such an assumption are analysed theoretically in Section 2.2. Moreover, the assumption is directly tested in the experiment (see Section 3 for the details).

2, and 3) extend to CARA preferences. We clarify the extent to which each hypothesis can be further generalized as we present it. Traders' preferences are represented by an individual utility function over wealth levels:

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

The solution of the trader's maximization problem under these assumptions yields (see, e.g., Gjerstad (2005)):

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

for the case of a risk-averse trader, 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 a risk-neutral or risk-loving one.

Under expected utility, and for any utility function, the optimal demand is zero when $p = b_i$; the optimal demand decreases as p increases if the utility function features non-increasing risk aversion (He and Treich (2017)). The optimal demand schedules for different levels of the CRRA coefficient are depicted in Figure 1. Risk-neutral and risk-loving players invest their entire endowment in long (short) positions whenever the price is below (above) their beliefs about the probability that the security pays 100. The demand of risk-averse traders is instead smooth. The amount invested is positively correlated with the distance between the price and the trader's belief. At any given price, the amount invested decreases as the degree of risk aversion increases. This observation establishes a first testable implication of the role played by risk aversion in the model.

HYPOTHESIS 1. *At the individual level, the net demand (in absolute terms) decreases as the degree of risk aversion increases.*

Hypothesis 1 extends to generic risk preferences provided that traders can be ranked in terms of their risk aversion. In fact, for any given price and belief, a lower demand reduces continuously the expected payoff as well as their variance. By inducing different demand schedules, risk aversion also affects market prices. To keep the presentation simple, let us now assume homogeneous risk preferences θ .⁸ We further assume that the

⁸The case of heterogeneous risk preferences is explored in Fountain and Harrison (2011) with a simulation exercise. In the empirical analysis, we take into account the individual heterogeneity within markets. See Section 5.




FIGURE 1. Optimal demand schedule and risk aversion. *Note:* The figure shows the optimal demand schedule of a trader holding a belief b_i corresponding to the dashed horizontal line for different degrees of risk aversion θ_i .

market as a whole does have information to share, imposing that average (and median) beliefs are informative, that is, $\bar{b} > 50$ ($\bar{b} < 50$) when $e = Blue$ ($e = Red$).

Log-utility ($\theta = 1$) sets an important benchmark, because in this case the market-clearing price coincides with the average belief of the traders: $p_{\theta=1}^* = \bar{b}$. $\theta = 1$ defines a proper prediction market (Wolfers and Zitzewitz (2006)), as one can directly infer the traders’ average estimate that an event will occur by observing the market price. In other words, under log-utility the price directly measures the amount of information in the market.

Lower degrees of risk aversion ($\theta < 1$) push prices toward the uninformed prior: $\bar{b} < p_{\theta < 1}^* < 50$ when $e = Red$, and $50 < p_{\theta < 1}^* < \bar{b}$ when $e = Blue$.⁹ As shown in Gjerstad (2005), this effect gets continuously stronger as risk aversion decreases and prices are the closest to the uniform prior when traders are risk-neutral or risk-loving (Manski (2006)). When risk aversion is higher than under log-utility ($\theta > 1$) the market price is instead closer to the true state of nature than the average belief: $p_{\theta > 1}^* < \bar{b} < 50$ when $e = Red$ and $p_{\theta > 1}^* > \bar{b} > 50$ when $e = Blue$.

Therefore, risk aversion shapes the equilibrium prices together with the amount (and distribution) of information in the market. Given the same fundamentals, prices are closer to the true state the more risk averse are the traders. The following testable implication summarizes this prediction, where p_{HIGH}^* (p_{LOW}^*) is the market-clearing price in a market with high (low) risk aversion.

⁹This pattern replicates the so-called favorite-longshot bias, an empirical regularity according to which unlikely states are overpriced, and likely states are underpriced (see, for instance, Snowberg and Wolfers (2010)).

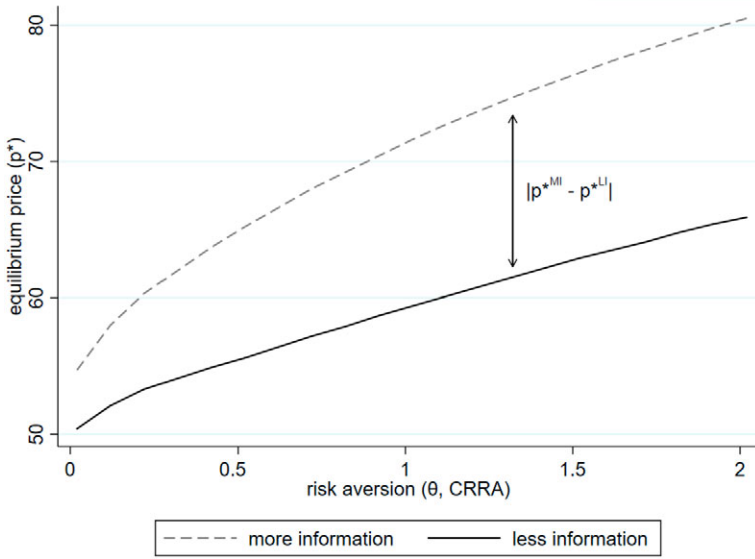


FIGURE 2. Visual representation of Hypotheses 2 and 3. *Note:* The figure shows the equilibrium prices for increasing degrees of risk aversion in markets characterized by different levels of information (MI vs. LI). The prices are derived computationally using the empirical distributions of beliefs implemented in our experiment. The distance of each of the two curves from 50 illustrates Hypothesis 2. The distance between MI and LI describes Hypothesis 3.

HYPOTHESIS 2. *Given the same amount of information, in more risk-averse markets the difference between market prices and the uninformed prior is higher than in less risk-averse markets:*

$$|p_{\text{HIGH}}^* - 50| > |p_{\text{LOW}}^* - 50|.$$

Figure 2 illustrates Hypothesis 2 in our experimental setting characterized by asymmetric beliefs. Note that for both levels of information the distance of the equilibrium price from 50 is increasing in the level of risk aversion.

Another testable implication can be derived exploiting the interaction between risk aversion and the amount of information. It can be shown that in more risk-averse markets the price should react more to the same increase in the amount of information. Consider one market in two different situations, with More Information (MI) and Less Information (LI) about the state, so that $\bar{b}^{\text{MI}} < \bar{b}^{\text{LI}}$ when $e = \text{Red}$ and $\bar{b}^{\text{LI}} < \bar{b}^{\text{MI}}$ when $e = \text{Blue}$. The extent to which market prices reflect the difference between the beliefs in the two situations also depends on θ . In particular, in the more risk-averse market prices react more to having more precise information, leading to the following hypothesis.

HYPOTHESIS 3. *The difference between market prices that reflect different amount of available information is higher in more risk-averse markets than in less risk-averse ones. That is, for each state:*

$$|p_{\text{HIGH}}^{\text{MI}} - p_{\text{HIGH}}^{\text{LI}}| > |p_{\text{LOW}}^{\text{MI}} - p_{\text{LOW}}^{\text{LI}}|.$$

Figure 2 also provides a visual representation of Hypothesis 3. The distance between the equilibrium prices increases with risk aversion when comparing markets with different levels of information.¹⁰

2.2 Informative prices

In a rational-expectation equilibrium, prices reveal the state to an observer that knows the equilibrium-price correspondence as long as the equilibrium prices differ in the two states. Traders fully trust the information contained in prices and disregard their private information. In this set-up, the no-trade theorem holds. Intuitively, by adjusting beliefs to the price, that is, setting $b_i = p$ in equation (2), the optimal demand is zero at any $0 < p < 100$. Since there exists no residual uncertainty in equilibrium, risk aversion does not play a role in this case.

In the prior information, equilibrium traders instead do not take this information into account at the time of submitting their limit orders. Although the rational expectation scenario seems unlikely in practice, also milder versions of strategic behavior create problems for the identification of Hypotheses 1–3 above. As an illustrative example, consider an individual that weights linearly his belief and the hypothetical price, interpreted as a market signal on the probability that $e = \textit{Blue}$:

$$b'_i = \alpha p + (1 - \alpha)b_i.$$

Figure 3 illustrates this situation. The dark dashed line represents the demand of a nonstrategic trader ($\alpha = 0$) characterized by $b_i = 50$ and $\theta_i = 1$. Consider now the case of $0 < \alpha < 1$. At $p = 50$, the information extracted from the market concurs with the trader's belief and, therefore, the choice of a strategic trader coincides with that of a nonstrategic player (point A). At a higher price, for example, $p = 70$, the market is signaling a probability that $e = \textit{Blue}$ higher than the trader believes. He readjusts his belief accordingly. The extent of the readjustment depends on α . Figure 3 shows the case of $\alpha = 0.5$. In this case, the trader behaves as if he believes that the probability that $e = \textit{Blue}$ is equal to 60% at $p = 70$: locally, his individual demand is derived from the gray line q_{60} , with point B_2 representing his net demand at $p = 70$. Similarly, at $p = 30$, his net demand corresponds to point C_2 . Iterating this procedure for all prices, gives the solid black line, representing the demand schedule of a strategic trader characterized by $b_i = 50$, $\theta_i = 1$, and $\alpha = 0.5$.

Partially internalizing the informative content of market-clearing prices reduces the quantity demanded (supplied) at any price $0 < p < 100$ that is below (above) one's belief.¹¹ That is, it produces a change in the individual demand that is similar to that of a

¹⁰CRRA or CARA specifications are not necessary conditions for Hypotheses 2 and 3. However, they cannot be proved to hold for generic risk preferences under our asymmetric belief distribution. For symmetric belief distributions, He and Treich (2017) show that their mechanics hold in general under nonincreasing risk aversion. It should also be noted that the results are specific to the Prior Information Equilibrium framework and the binary asset structure. For instance, in a CARA-normal noisy rational expectations equilibrium, the information received by less risk-averse traders is overweighted in market prices, because these traders trade more aggressively (Hellwig (1980)).

¹¹The only exception occurs when the no-bankruptcy constraint is reached for different levels of α , which is always the case for risk-neutral and risk-loving traders.

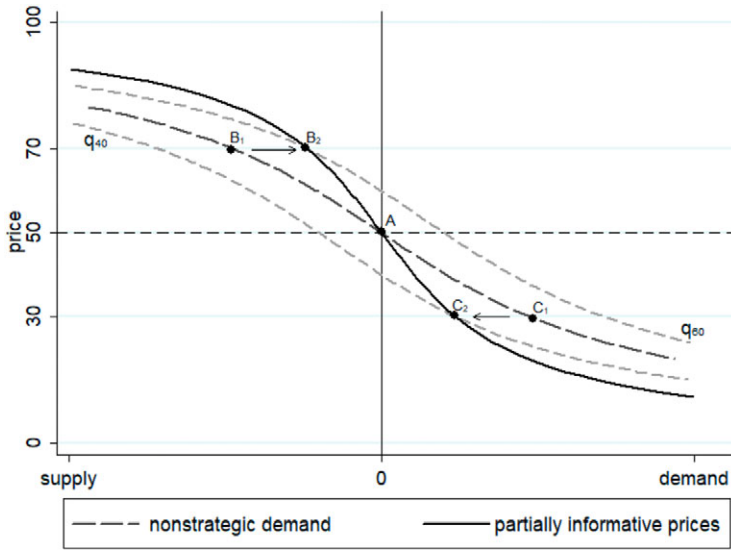


FIGURE 3. Optimal demand schedule with informative prices. *Note:* The figure shows the optimal demand of a trader ($b_i = 50$, $\theta_i = 1$) that takes (partially) into consideration the information contained in market-clearing prices. If he did not, his demand would be the dark dashed line. The solid dark line is his demand if he anticipates the information contained in prices, by updating his beliefs according to $b'_i = 0.5p + 0.5b_i$. This demand is derived, at any given price, as the optimal demand of a trader whose beliefs are halfway between the price and his original beliefs (e.g., the light gray lines).

higher degree of risk aversion (see Figure 1). Since an increase of θ_i may be confounded with a larger α , the assumption of nonstrategic behavior is crucial for identifying the role played by risk aversion. Acknowledging this problem, we design a specific treatment to directly test for the validity of the prior information assumption corresponding to $\alpha = 0$.

3. DESIGN

Elicitation of risk preferences

At the beginning of the experiment, we elicit an independent measure of 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 Monetary Units (MU) between a safe account and a risky investment. The latter 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 (Crosetto and Filippin (2016)). We label θ_i^{inv} the CRRA coefficients assigned to participants from their choice in the Investment Game. As shown in Section 2, log-utility has special features in the model because it implies that $p^* = \bar{b}$. The Investment Game is superior to other tasks in scan-

ning risk preferences around this level.¹² The variation in individual risk preferences, as measured by θ_i^{inv} , allows us to test for Hypothesis 1.

Matching

We divide each session in two groups of 11 traders, according to their choice in the Investment Game, separating the subjects above and below the median of the session. 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 each market. This manipulation allows us to test for Hypothesis 2.

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 they know that each urn is selected with equal probability. It is common knowledge that all subjects start with the same uniform prior over urns.

A simple asset called “Majority Blue” is traded in the market. If the urn is C or D, the state “the majority of marbles are Blue” realizes ($e = \text{Blue}$) and every asset pays the owner 100 MU at the end of the trading period. If the urn is A or B, the state “the majority of marbles are Blue” does not realize ($e = \text{Red}$) and every asset pays the owner 0 MU.¹³

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, that is, $s \in \{x - 5, \dots, x + 5\}$ 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 distributes the signals is also common knowledge.

Given the signal, Bayesian updating of the prior about the state generates the posterior beliefs $p(e = \text{Blue}|s)$ reported in the last row of Table 1. Some signals ($s \leq 45$ and $s \geq 55$) reveal the state with certainty, leading to $p(e = \text{Blue}|s) \in \{0, 1\}$. Other signals ($s = 46, 47, 53, 54$) are partially informative, that is, $p(e = \text{Blue}|s) \in \{\frac{1}{3}, \frac{2}{3}\}$. Finally, some signals ($48 \leq s \leq 52$) are uninformative and, therefore, $p(e = \text{Blue}|s) = \frac{1}{2}$. Aggregating $p(e = \text{Blue}|s)$ over subjects is straightforward to compute that the average beliefs about

¹²The Holt and Laury task (Holt and Laury (2002)) and the Bomb Risk Elicitation Task (Crosetto and Filipin (2013)) include log-utility only as a limit case. On the other hand, it is not a concern that the Investment Game does not discriminate parameters weakly lower than 0, while the other tasks can, because the optimal behavior of risk-neutral and risk-loving traders does not differ, according to the model, also in the market.

¹³Note that the urn selected is deterministically linked to the value of the asset; there is no draw from the urn.

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

Note: 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.

the state $e = Blue$ are 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.

The combination of the underlying urn with the structure of the signal constitutes a within-subject manipulation of the amount of information in the market. In fact, urns *A* and *D* provide more information than urns *B* and *C* on the realization of each of the two states. In conjunction with the matching protocol, this manipulation allows us to test for Hypothesis 3.

Belief elicitation

Prices depend on traders’ beliefs and, therefore, not updating the beliefs in a Bayesian manner would greatly affect price formation. Furthermore, our tests of hypotheses would be confounded if failures of Bayesian updating correlate with risk preferences. 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 (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 , that is, if the subject’s beliefs are sufficiently accurate, he earns a fixed prize (200 MU); otherwise he gets nothing. The BSR 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 different probabilities 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.¹⁴

¹⁴The QSR 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 utility of reducing the variance of the outcomes more than compensates the lower expected reward.

Market institution

The market institution is a CA. The choice of the CA is motivated by the fact that identifying the role played by risk aversion in the information aggregation process relies on the prior information assumption. This assumption is more likely to be violated under different market institutions that involve real-time strategic interaction among the traders.

We implement the CA as follows. Subjects have 2 minutes in which they can place limit bid and ask orders for the asset. Limit orders are combined to form net demands. A trader's demand at a price p is the sum of the orders that would be executed at that price, that is, buy orders with a limit price equal or higher than p , and sell orders with a limit price equal or lower than p . As they submit orders, a visual representation of their demand schedule updates in real time. At the end of the 2 minutes, the equilibrium price is computed as the price that equalizes aggregate demand and supply, maximizing the volume of trades. All of the orders with compatible limit prices are executed at that price.¹⁵

In order to properly investigate the role of risk aversion in information aggregation, it is important to shut down other potential channels through which risk aversion can play a role. A setting without aggregate risk avoids to impute an a priori bias in the price formation, as stressed by [Bossaerts, Frydman, and Ledyard \(2013\)](#). Given that the asset has common value, making the endowment orthogonal to the realized state ensures no aggregate risk. Therefore, subjects enter the market with 1000 MU and no asset in each trading period. This feature also prevents the interaction between risk preferences and the evaluation of a risky endowment, which may influence behavior in complicated ways. 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 that are net sellers at the market-clearing price 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 pending orders, making sure that the net demand does not require more than the endowment for any possible market price.

The CA represents the ideal test-bed to study the role of risk aversion on information aggregation because it provides well-defined individual net demands. Furthermore, the demand schedules submitted in the CA allow us to estimate an individual parameter of risk aversion. A natural exercise is therefore to test whether risk preferences are consistent across contexts.¹⁶

HYPOTHESIS 4. *The individual coefficient of risk aversion estimated from the market behavior is consistent with the measure elicited with the Investment Game.*

¹⁵In 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.

¹⁶Note that a rank-based test of Hypothesis 4 does not rely upon CRRA specification. As long as individual demands can be ordered in terms of risk aversion, the rank order of the coefficients does not depend on the specific functional form.

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	0.01	0.09	0.20	0.40	0.20	0.09	0.01

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

Random price mechanism

As discussed in Section 2.2, subjects may extract information from the market prices through hypothetical reasoning, also in the CA. This possibility would call the prior information equilibrium into question, and confound the identification of how risk aversion shapes information aggregation. Therefore, we directly test for the validity of the prior information assumption by running an additional treatment in a between-subject design. In this treatment, we manipulate the price formation mechanism *ceteris paribus* implementing a Random Price Mechanism (RPM) á la Becker, DeGroot, and Marschak (1964). At the end of the order-submission phase, one price is randomly drawn for each trader, from a known distribution. The probability distribution of the prices mimics the empirical frequency of prices observed in the CA treatment, and is represented in Table 2.¹⁷ Since it is common knowledge that this distribution does not depend on the selected urn, prices do not carry any informative content, and subjects cannot behave strategically under the RPM. Under the prior information assumption, the behavior of traders does not differ between the CA and the RPM.

HYPOTHESIS 5. *The demand schedules do not differ significantly between the call auction and the random price mechanism.*

4. PROCEDURES AND PAYMENT SCHEME

The sessions were run between February and September 2018 at the Experimental Economics lab (EELAB) of 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. Only in the CA treatment, the software ranks the subjects' choices and assigns them to one of two markets of 11 traders (high or low risk aversion).

To prevent compound-risks considerations when choosing in the Investment Game, subjects receive detailed instructions on the rules and the working of the market after they made their decision. During the instructions, they are asked to answer a battery of quizzes that assess their comprehension of the various parts of the instructions.¹⁸

¹⁷The reason why prices are not drawn from a uniform distribution is that drawing more often prices in the tails of the support as compared to the CA treatment could affect the choices regardless of the strategic behavior of the subjects.

¹⁸The 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

The reading of the instructions would move on only once all subjects have cleared the quiz. For each quiz, we keep track of the number of mistakes each subject makes before clearing it.

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 order submission phase, subjects are informed of the price in that period, their liquidity and asset portfolio, but not of the selected urn. Then they are asked to report again the probability that each urn has been selected.

In order to avoid spill-over effects from the outcomes of previous periods or tasks, all uncertainty is resolved at the end of the experiment. The computer selects randomly: (i) the outcome of the Investment Game for each participant; (ii) one period for all participants to be used for payments of the trading task; (iii) one of the measures of beliefs for each participant in a period that is not relevant for the payment of the trading task (preventing hedging between beliefs and market behavior). 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 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. Subjects are then notified about their earnings, they fill in a short questionnaire, and finally they receive their compensation in an anonymous manner.

All participants are undergraduate students recruited through the lab recruitment platform. In the CA treatment, we collect data from 10 sessions, or 220 experimental subjects. There are therefore 20 independent observations—10 high- and 10 low-risk aversion markets—with data on 240 trading periods, 60 for each urn. In the RPM treatment, we collect data from 38 experimental subjects, divided into 2 sessions. Given the rules of the RPM, each subject represents one independent observation. Sessions lasted about 2 hours and the average payment was 16.20€.

5. RESULTS

We open the section with the comparison of treatments CA and RPM, in order to validate the prior information assumption. In the following subsections, we analyze how elicited risk aversion correlates with trading activity at the individual level (Hypothesis 1) and with market outcomes (Hypotheses 2 and 3), respectively. We then use individual demands to estimate subjects' risk aversion directly from their market activity. The behavior implied by such estimates is at odds with the observed pattern of prices, thereby inducing us to enrich the specification of the individual demands. Incorporating

working of the market interface, and also includes 2 minutes to interact freely with the interface. Complete instructions are attached in the A.

the tendency to use only in part the information possessed allows us to rationalize the market outcomes. Finally, we look at how the estimated parameters, both representing risk aversion, correlate with the choices in the Investment Game, thereby testing Hypothesis 4. As common practice in the experimental literature, equality of distributions across populations are tested nonparametrically. The reason is that parametric assumptions on the shape of distributions are usually violated in small samples.¹⁹

5.1 Test of the prior information model

According to Hypothesis 5, we test the prior information assumption by comparing whether behavior significantly differs in treatment CA and RPM. Subjects anticipating the informative content of prices would result, for any price, in a reduced absolute net demand in CA with respect to RPM, which for the purpose of the graphical comparison, one can simply think of as “steeper.” Such an effect is not visible in Figure 4, which displays the average net demand in the two treatments over all periods. We compare the distribution of the average absolute quantities demanded at each price between CA and RPM using a Kolmogorov–Smirnov test. The test fails to detect a significant difference ($KS = 0.1475$, $p = 0.520$). Alternative testing strategies exploit the fact that strategic

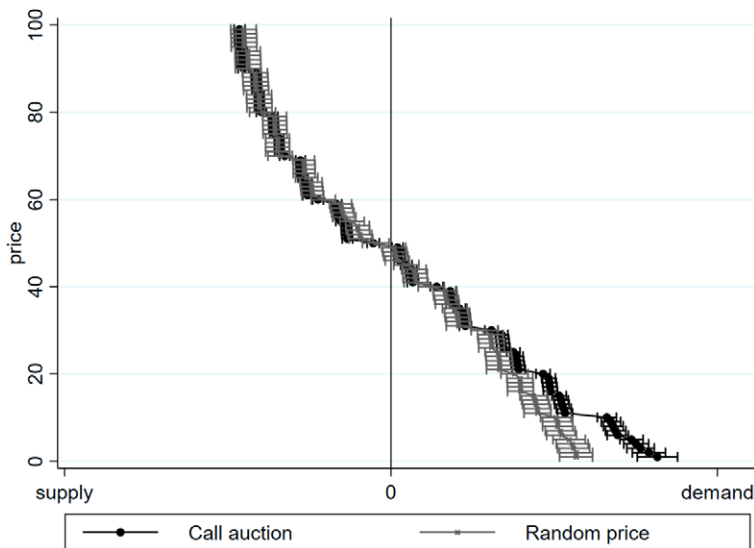


FIGURE 4. Average demand in treatment CA and RPM. *Note:* The figure shows the average aggregate net demand (over all periods and urns), for the call auction markets and the random price mechanism, with 95% confidence intervals for each price.

¹⁹The most common is the Mann–Whitney U test, which tests the null hypothesis that, for randomly selected values X and Y from two independent samples, the probability of $X > Y$ is equal to the probability of $Y > X$. The Wilcoxon signed-rank test is its counterpart for the case of two dependent samples (e.g., urn A vs. urn B for the same market). The Kolmogorov–Smirnov test is instead based on the distance between two observed cumulative distributions. In all cases, the statistics are computed without imposing a specific shape of the underlying distributions.

behavior would be confounded with a higher degree of estimated risk aversion. We estimate the CRRA coefficients from the individual demand schedules, as well as the slopes of linear approximation of the demand schedules. Comparing either of these estimates across CA and RPM with Mann–Whitney U tests, using one observation per subject, we always fail to reject the null of no difference.²⁰ Since the individual behavior does not differ across conditions, we conclude that the prior information model passes the test and in what follows we restrict the analysis to the data from the CA treatment.

5.2 Elicited risk aversion: Individual choices

In the Investment Game, subjects invest on average 93.5 out of 200 ECU, while the median choice is 100. The corresponding average CRRA coefficient is $\bar{\theta}^{\text{inv}} = 0.71$ (median 0.32). The median seems to signal a relatively low degree of risk aversion but this low value as well as its discrepancy with the mean depend on the inner working of the task.²¹ Indeed, decisions in our experiment are far from the implication of risk neutrality, which requires that the whole endowment is invested: only 10% of the subjects opt for such a choice.²²

The theoretical model described in Section 2 posits that risk aversion negatively correlates with the quantities demanded or supplied at the individual level. In fact, the higher the degree of risk aversion, the lower the trading activity at any price. Our experimental results confirm that this is the case. The individual degree of risk aversion θ_i^{inv} negatively and significantly correlates with the average number of assets demanded/supplied (Spearman's $\rho = -0.24$, $p < 0.01$). A subject in the bottom quartile of the distribution of θ_i^{inv} (i.e., the least risk averse) has a net demand, that is, on average, 27% larger than that of a subject in the top quartile.

RESULT 1. The degree of risk aversion negatively correlates with the quantities demanded and supplied at the individual level.

5.3 Elicited risk aversion and market prices

Hypotheses 2 and 3 summarize the predicted effect of risk aversion on information aggregation, that is, market-clearing prices. The two hypotheses are derived under the assumption of homogeneous risk preferences. Coherently, our protocol assigns participants to markets trying to minimize the heterogeneity of risk preferences within mar-

²⁰CRRA coefficients: $U = -0.612$, $p = 0.541$; Slopes: $U = -0.532$, $p = 0.595$. The estimation of CRRA coefficients from the demands represents the core of Section 5.4. We report the details of the estimation procedure in that section.

²¹The transformation of choices into θ_i^{inv} coefficients is highly nonlinear (see Figure 3 in Crosetto and Filippin (2016)). θ_i^{inv} is more sensitive to changes in the choice for low amounts invested; a low level of risk aversion ($\theta_i^{\text{inv}} = 0.32$) already emerges when half of the endowment is kept; θ_i^{inv} decreases then very slowly for more risk-taking decisions. For this reason, in the analysis we will often rely upon Spearman, rather than linear, correlations.

²²The cumulative distribution of all the choices is reported in Figure C.1 of Appendix C.




FIGURE 5. Market prices in low- and high-risk aversion markets. *Note:* The boxplot shows the distribution of market prices for each urn, separately for high and low risk-aversion markets (“high RA” and “low RA,” respectively).

kets. At the same time, it maximizes the variance across markets in order to test the role played by a different level of risk aversion.²³

Figure 5 summarizes the distribution of prices in both high and low risk-aversion markets, given the selected urn. Prices contain relevant information, because they significantly differ between any pair of urns according to a battery of Wilcoxon signed-rank test, both in high and in low risk-aversion markets (see the top panel of Table 3).²⁴

The central panel of Table 3 compares equilibrium prices in high and low risk-aversion markets by urn. Hypothesis 2 predicts that prices should be lower in high risk-aversion than in low risk-aversion markets when the urn is *A* or *B*, and vice versa when the urn is *C* or *D*. A Mann–Whitney *U* test fails to detect a significant difference for urns *A*, *B*, and *C*. Where a significant difference emerges, that is, in urn *D*, it goes in the opposite direction than predicted by Hypothesis 2.

RESULT 2. The distance between market prices and the uninformed prior is not higher in high risk-aversion markets.

²³The results presented in this section do not depend on the residual variance of risk preferences within markets. On the one hand, the variance within markets is of second-order importance as compared to the variance across markets. The average choice in the Investment Game ranges between 40 and 69 in high risk-aversion markets, and between 119 and 150 in low risk-aversion markets. On the other hand, Figure C.2 in Appendix C shows the predicted market-clearing prices by urn given the joint distribution of observed signals and elicited risk preferences. Our hypotheses survive to the heterogeneity of risk preferences in our markets.

²⁴All nonparametric tests on market outcomes are intended as using one independent observation per market.

TABLE 3. Nonparametric tests of hypotheses.

	Urns <i>A</i> vs <i>B</i>		Urn <i>B</i> vs <i>C</i>		Urn <i>C</i> vs <i>D</i>	
	<i>z</i>	<i>p</i> -value	<i>z</i>	<i>p</i> -value	<i>z</i>	<i>p</i> -value
All mkts	-2.315	0.021	-3.155	0.002	-2.895	0.004
Low	-1.070	0.285	-2.756	0.006	-2.090	0.037
High	-2.191	0.028	-1.682	0.093	-2.095	0.036

	Urn <i>A</i>		Urn <i>B</i>		Urn <i>C</i>		Urn <i>D</i>	
	<i>U</i>	<i>p</i> -value	<i>U</i>	<i>p</i> -value	<i>U</i>	<i>p</i> -value	<i>U</i>	<i>p</i> -value
Low vs. High	-0.907	0.364	-0.416	0.677	-1.665	0.096	-2.307	0.021

	Urn <i>A</i> – Urn <i>B</i>		Urn <i>D</i> – Urn <i>C</i>	
	<i>U</i>	<i>p</i> -value	<i>U</i>	<i>p</i> -value
Low vs. High	-0.303	0.762	-0.227	0.820

Note: The top panel reports the Wilcoxon signed-rank test, and corresponding *p*-value, on the difference in prices between pairs of urns, both on aggregate and separately for low and high risk-aversion markets. The central panel reports, for each urn, the Mann–Whitney *U* test statistic, and corresponding *p*-value, on the difference in prices between low and high risk-aversion markets. A positive statistic means a higher value for high risk-aversion markets. The bottom panel reports the Mann–Whitney *U* test statistic, and corresponding *p*-value, on the difference between low and high risk-aversion markets in the difference in prices between the pairs of urns *A/B* and *C/D*. A positive statistic means a higher value for high risk-aversion markets. All statistics are computed using one observation per market (20 independent observations, 10 low and 10 high risk aversion).

Despite the low number of independent observations in absolute terms (20), this result unlikely represents a false negative. In fact, given the average CRRA coefficient observed in the high and low risk-aversion markets, 1.19 and 0.23, respectively, the expected difference in market prices is substantial: 7 ECU in urns *B* and *C*, and 13 ECU in urns *A* and *D*. Given the observed standard deviations of market prices, we have a power of 0.8 (with a confidence level of 0.95) to detect a difference in prices of about 9.2 ECU. Our experiment is therefore sufficiently powered to detect the expected effect size in high information markets, and has a power of 0.61 in the low information markets.

The average CRRA coefficient in the high risk-aversion markets is rather close to log utility. However, prices by urn are not close to the average beliefs (28.8, 40.9, 59.1, 71.2) as the model predicts. Fully exploiting the variance of θ_i^{inv} at the market level confirms that the level of risk aversion does not shape the equilibrium price. In more detail, Figure 6 displays the absolute distance of the observed equilibrium prices from the average Bayesian belief, and its relation with the average degree of risk aversion ($\bar{\theta}^{inv}$) in each market. The overimposed line is a linear fit showing that the two measures are virtually orthogonal, while according to the model the relationship should be downward sloping and cross the horizontal axes when $\theta^{inv} = 1$.²⁵

Hypothesis 3 predicts that prices are more responsive to the amount of information distributed the higher the degree of risk aversion in the market. For each market, we

²⁵This conclusion is robust to using the average elicited belief, rather than the Bayesian one, as shown by Figure C.3 in Appendix C. In other words, this result cannot be explained by failures of Bayesian updating.

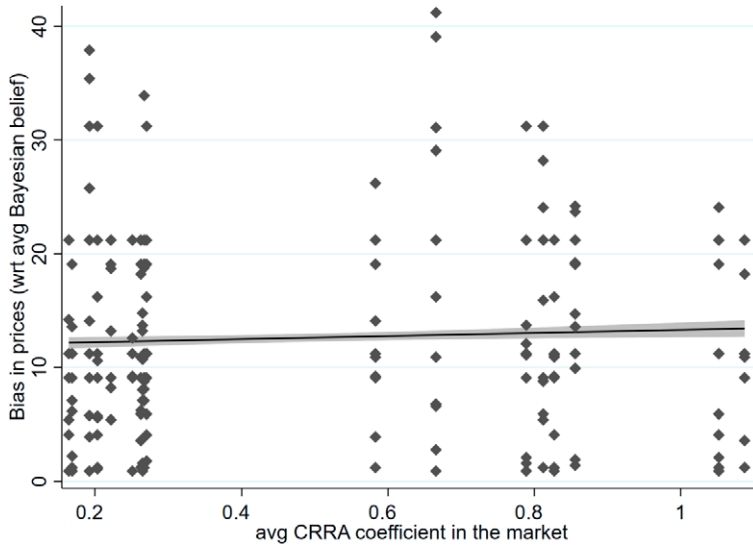


FIGURE 6. Distance of price from average belief and risk aversion. *Note:* The figure plots the distance of market prices in each market/period from the average Bayesian belief, plotted against the average CRRA coefficient in the market (the gap between 0.27 and 0.58 is the effect of the matching procedure). A linear fit between the two measures is overlaid (with 95% CI based on the standard error of the prediction).

compute the average difference in prices between urns A and B , and between urns C and D , that is, for each state we measure how prices react to a larger amount of information. We then test for differences in these differences between high and low risk-aversion markets. Results are reported in the bottom panel of Table 3, showing an indistinguishable pattern. We conclude therefore against Hypothesis 3.

RESULT 3. Market prices in high risk-aversion markets are not more responsive than in low risk-aversion markets to the amount of information distributed.

Risk neutrality has long been the standard assumption for the prior information equilibrium in experiments on information aggregation in markets (Plott and Sunder (1982, 1988), Choo, Kaplan, and Zultan (2017)). Since in our data the risk aversion of traders does not affect prices, one could argue that such evidence corroborates the risk neutrality assumption in deriving the predicted equilibrium prices. 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 we have seen that 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 .

While at first glance our results seem to speak in favor of risk neutrality, different indicators suggest 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. Moreover, we have shown that subjects with a higher θ_i^{inv} tend to have a lower exposure in the market (Result 1). In the next section, we delve deeper into this conundrum by estimating individual risk preferences from trading behavior and showing that inferring risk neutrality from the observed prices would indeed be a misleading exercise.

5.4 Estimate of individual risk aversion

In this section, we estimate CRRA coefficients from the individual demand schedules in order to understand how choices of risk averse traders result in roughly risk-neutral prices. To do so, we first consider equation (2), imputing the actual endowment $m = 1000$ and imposing the Bayesian beliefs ($0 < b_i < 100$) implied by the signal received by each subject in every period. We then use nonlinear least squares to estimate the individual CRRA coefficient that best explains the observed individual net demand (q_i) at any price $20 \leq p \leq 80$. We label this coefficient θ_i^{mkt} , to distinguish it from θ_i^{inv} .²⁶ The reason to limit the price range is that subjects trade systematically less than predicted at extreme values of the prices, as shown by Figure 4, likely because they do not bother placing orders at prices they deem implausible.²⁷

The average estimated coefficient is $\theta^{\text{mkt}} = 2.92$. Even relying upon the median, which is not sensitive to outlier decisions, we find a value (1.86) of a different order of magnitude with respect to the elicited θ_i^{inv} (mean = 0.71, median = 0.32). Consequently, the prices predicted according to the estimated θ_i^{mkt} are even farther from the observed ones than those computed using θ_i^{inv} (see Figure 7). 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.

Figure 8 shows the net demand by urn aggregated across all markets, together with the corresponding predicted one, obtained aggregating individual demands computed using the estimated θ_i^{mkt} . In the theoretical model, risk aversion explains the shape of the net demand, that is, how much subjects are willing to trade as the price moves away from one's beliefs. In the range of prices between 20 and 80, demands are close to linear and, with a semantic abuse, we will from now on refer to θ_i^{mkt} as the slope of the demand because it is visually more intuitive. The intercept represents an estimation of the market-clearing price. Figure 8 shows that the estimated values of θ_i^{mkt} correctly predict the slope of the aggregate net demands. The problem is the *intercept* of the demands, which is higher than predicted for urns *A* and *B*, and lower for urns *C* and *D*. The observed aggregate demands always cross the vertical line (and consequently identify an

²⁶In principle, $\theta_i^{\text{mkt}} \rightarrow \infty$ for a subject that never trades. For the sake of comparability, we cap the maximum value of θ_i^{mkt} to 32.48, which is the maximum level attributable within the Investment Game.

²⁷Indeed, only in 6 out of 240 cases the market price falls below 20 or above 80. Including also $p < 20$ and $p > 80$ would overestimate the parameter for the most relevant range of prices, but results below are qualitatively robust even using the full support of prices and are available upon request. The same holds for Result 1 ($U = -1.097$, $p = 0.273$).

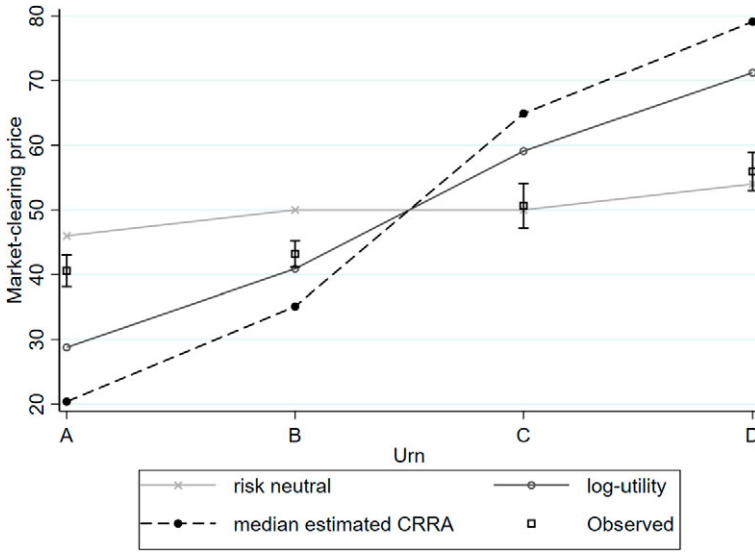


FIGURE 7. Observed average prices and predicted prices for various CRRA coefficients. *Notes:* The figure shows the predicted price for each urn in a risk-neutral market, in a log-utility market, and in a market where all traders have the median CRRA coefficient estimated for each subject from his demand schedule. The squares are the average observed prices over all markets (with 95% confidence intervals).

equilibrium price) closer to 50. This explains how high estimated risk-aversion coefficients can coexist with prices close to the risk-neutral benchmark.

The underlying mechanism is better illustrated without aggregating over different signals. Figure 9 displays the average net demand of all subjects when receiving a signal $s \in \{53; 54\}$, which implies $b_i = 66.6$ but the same argument applies in our data to all the partly or fully informative signals. According to the theory, traders should switch from buyer to seller around their Bayesian belief. Figure 9 shows that subjects switch instead at a lower price, around 55. In other words, they start selling short below their expected value of the asset.²⁸

An obvious candidate to rationalize this finding is conservatism, proposed by Peterson and Miller (1965) who posit that variations in subjective probability revision are in the same direction, but of smaller magnitude, than corresponding variations in Bayesian probability change. Our experimental protocol allows us to scrutinize the role of failures of Bayesian updating. Table 4 reports the subjects' beliefs about the state $e = Blue$ before the market opens, for all participants and separately for the top and bottom half of the distribution of estimated risk-aversion coefficients θ_i^{inv} . When receiving fully informative signals, subjects update almost perfectly, and declare they are sure about the state more than 90% of the times. When receiving partially informative signals subjects instead show evidence of overreaction, that is, the opposite of conser-

²⁸Note that the net demand displays some price insensitivity around zero like in Fattinger (2018), but the whole interval of prices for which the demand stays close to zero lies well below the Bayesian belief.

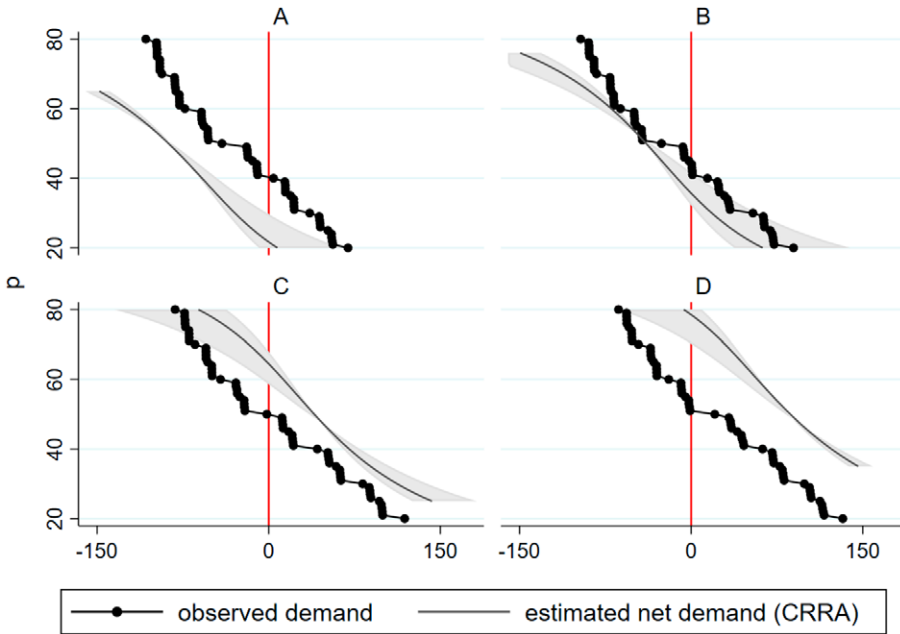


FIGURE 8. Observed and predicted aggregate net demands. *Notes:* The figure shows for each urn (i) the average observed aggregate demand, and (ii) the average aggregate predicted demand (with 95% confidence interval). This is obtained aggregating individual optimal demands, obtained using the CRRA coefficients estimated for each individual from his demand schedules (θ_i^{mkt}).

vatism. Therefore, conservatism in the way it has been defined cannot account for the observed demands.²⁹ Moreover, the magnitude of errors with respect to Bayesian updating does not correlate with risk aversion as estimated from the Investment game (Spearman’s $\rho = 0.07$, $p = 0.294$), nor do errors or beliefs differ significantly between low and high risk-aversion participants (Errors: Mann–Whitney $U = -0.764$, $p = 0.445$; Beliefs: Mann–Whitney $U = -0.635$, $p = 0.526$).

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 in beliefs, is transferred to the choices only in part. Therefore, this mechanism refers to actions rather than to beliefs. The transformation of subjective probabilities toward 50:50 has been documented in the literature, particularly in complex and ambiguous environments (Asparouhova, Bossaerts, Eguia, and Zame (2015), Fattinger (2018)). Dimmock, Kouwenberg, and Wakker (2016) and Trautmann and Van De Kuilen (2015) label this inclination “a-insensitivity.” Enke and Graeber (2019) embed this attitude in an encompassing framework of cognitive uncertainty. Whatever the ultimate cause, a trader who uses his information only in part is behaviorally indistinguishable from

²⁹Average predicted demands obtained using elicited beliefs are not distinguishable from those depicted in Figure 8.

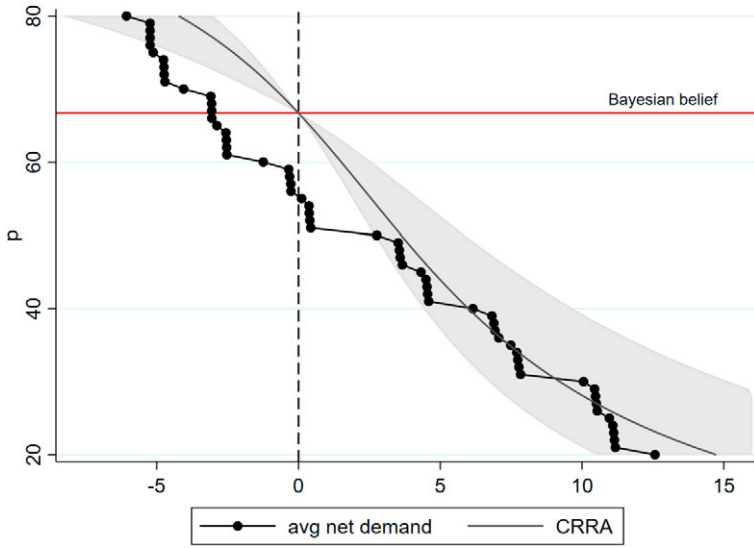


FIGURE 9. Observed and predicted aggregate net demand for $b_i = 66$. *Notes:* The figure shows (i) the average observed individual demand of subjects that received signals $s_i = 53, 54$ (implying a Bayesian belief of $2/3$), and (ii) the average predicted individual demand of the same subjects (with 95% confidence interval). This is obtained averaging individual optimal demands, obtained using the CRRRA coefficients estimated for each individual from his demand schedules (θ_i^{mkt}).

another trader who partially updates his subjective probability and then acts fully exploiting such conservative beliefs, and we label this attitude *operational conservatism*.³⁰

TABLE 4. Ex ante beliefs about the state $e = Blue$.

Signal received	Bayesian beliefs	Reported beliefs		
		Low RA	High RA	All
42–45	0	2.8	4.8	3.8
46–47	33.3	26.1	28.5	27.3
48–52	50	52.0	52.9	52.4
53–54	66.6	75.2	76.0	75.6
55–58	100	97.3	96.9	97.1

Note: The table displays the average self-reported beliefs about $e = Blue$ elicited from subjects before the market opens. This subjective probability is computed summing up the probability assigned to urns C and D for each subject. Each row represents a set of signals (left column) corresponding to a unique Bayesian posterior (second column). For each set of signals, the last three columns display the average self-reported beliefs, separately for low and high risk-aversion (RA) participants, as well as for the whole sample (All).

³⁰A similar inclination not to use the information possessed has been found by Fryer (2013) in an experiment on education and outcomes. Page and Siemroth (2021) model how much of the available information is used by mapping the observed prices onto well-defined subsets of available information. Their exercise cannot be replicated in our context because of an intrinsically different structure of the signals.

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 prior ($b_i = 50$):

$$\hat{b}_i = (1 - \delta_i)b_i + \delta_i(50). \tag{4}$$

We then reestimate the individual demand (equation (2)) substituting b_i with \hat{b}_i . The new specification includes the parameter δ_i meant to capture the amount of information that is not incorporated into the choices. The effect of δ is that of shifting the net demand, with the intercept moving toward 50 as $\delta_i \rightarrow 1$.

The two-parameter $(\theta_i^{\text{mkt}}, \delta_i)$ estimate of the individual demands reveals that a lot of information is not used. δ_i is significantly larger than zero for 141 subjects out of 219, with $\bar{\delta} = 0.64$. The median estimated value of θ_i^{mkt} decreases from 1.86 to 1.79 when δ_i is included in the model. 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_i^{mkt} and δ that we discuss in more detail in Section 5.5.

As shown by Figure 10, the $(\theta_i^{\text{mkt}}, \delta_i)$ specification allows us to properly reconstruct the aggregate behavior in the market. The gap between observed and predicted de-

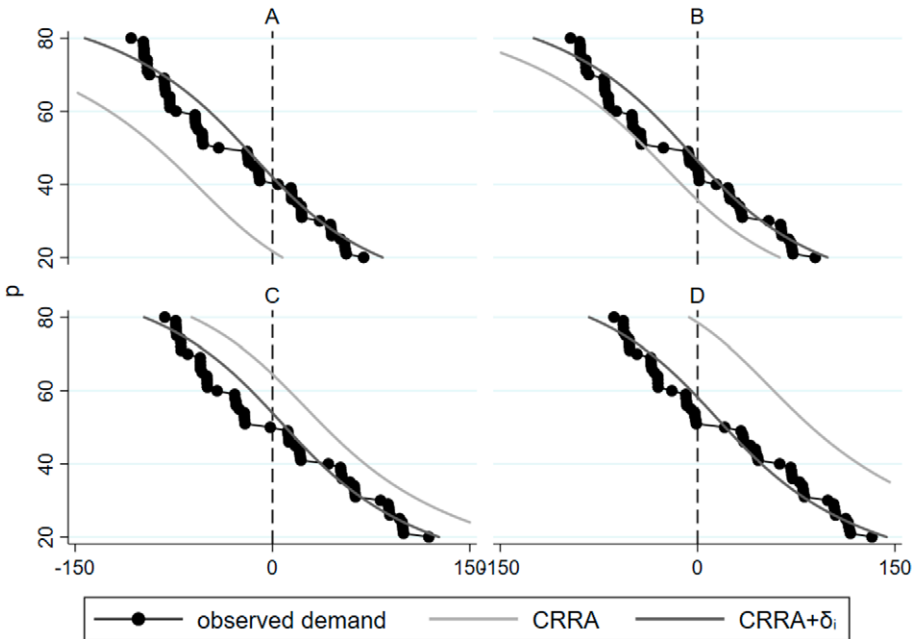


FIGURE 10. Observed and predicted (with and without δ) aggregate demands. *Note:* The figure shows, for each urn (i) the average observed aggregate demand, (ii) the average predicted demand obtained aggregating individual optimal demands computed using the estimated CRRA coefficients θ_i^{mkt} , and (iii) the average predicted demand obtained aggregating individual optimal demands computed using the estimated CRRA coefficients and operational conservatism (θ_i^{mkt} and δ_i).

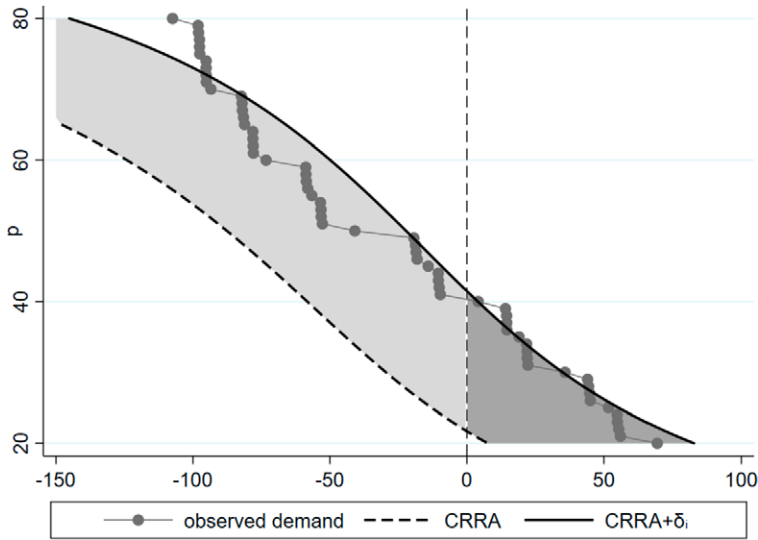


FIGURE 11. Operational conservatism and market exposure. *Note:* The figure shows, for urn *A*, (i) the average observed aggregate demand, (ii) the average predicted demand obtained aggregating individual optimal demands, obtained using the estimated CRRRA coefficients θ_i^{mkt} , (iii) the average predicted demand obtained aggregating individual optimal demands, obtained using the estimated CRRRA coefficients and operational conservatism (θ_i^{mkt} and δ_i). The dark (light) shaded area represent the increase (reduction) in market exposure due to operational conservatism.

mands disappears almost completely when the model accounts for the partial use the subjects make of the information they have.

An interesting feature of δ_i is that it induces a lower exposure in the market. Figure 11 displays the estimated aggregate demand, with and without δ_i in the model, when the underlying urn is *A*.³¹ Positive values of δ_i imply an aggregate demand that crosses the vertical axis around 40, that is, closer to 50 than what predicted by the estimated θ_i^{mkt} alone. The upward shift of the demand implies overall 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 $\delta_i > 0$ is therefore a less intense trading activity. Therefore, it could be interpreted as behaviorally related to risk attitudes.

Data at the individual level show on the one hand that the standard deviation of potential earnings in all the 12 periods negatively and significantly correlates with δ_i ($\rho = -0.136$, $p = 0.045$). On the other hand, the waste of information implies a suboptimal behavior and in fact also the average potential earnings negatively correlate with δ_i ($\rho = -0.50$, $p \leq 0.001$). In other words, $\delta_i > 0$ reduces the variance of the outcomes at a (inefficiently high) cost. In this sense, it may represent an additional behavioral manifestation of a subject's risk preferences.

One may argue that *operational conservatism* is a short-run phenomenon, which disappears as the subjects gain experience. We model a nonlinear learning process in the

³¹An identical argument applies to the other urns as well.

estimation of δ . In particular, we allow δ to change over the market periods according to a linear and a quadratic term. We run the estimation at the market level, and clustering errors at the individual level, so as to have enough observations. We indeed find that, on average, $\delta \approx 0.73$ in the first period and decreases over time. The model predicts that the speed of the learning process also decreases over time: it reaches $\delta \approx 0.58$ after half of the periods, and reaches a minimum of $\delta \approx 0.54$ in the last three periods.

The identification of *operational conservatism* allows us rationalize why risk averse traders may generate seemingly risk neutral prices. Even when the distribution of information in the market is fully controlled, we cannot take for granted that it will be fully exploited. In our experiment, the observed pattern of prices reflects the aggregation of a lower amount of information than that possessed by traders, rather than risk neutrality. An important corollary of this exercise is to show why trying to infer the average degree of risk aversion from market prices can be a misleading exercise.

5.5 Consistency of risk-aversion measures

As we have seen, CRRA coefficients estimated from individual demands are substantially larger than those implied by choices in the Investment Game. A possible interpretation for such a pronounced difference is that subjects perceive the market as a more risky environment. The virtually infinite set of outcomes without objective probabilities implied by the call market may induce 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. We find a significant Spearman's rank correlation between θ_i^{mkt} and θ_i^{inv} (Spearman's $\rho = 0.16$, $p = 0.015$).³²

It is interesting to verify whether δ_i correlates with the measure of risk aversion elicited independently. Indeed, θ_i^{inv} and δ_i are significantly correlated (Spearman's $\rho = 0.16$, $p = 0.020$). The lower the choice in the Investment Game, that is, the higher θ_i^{inv} , the higher δ_i , that is, the more operationally conservative is the subject. Table 5 shows that the choice in the Investment Game significantly correlates with θ_i^{mkt} and δ_i both separately (Columns 1 and 2) and at the same time (Column 3).

Since holding $\delta_i > 0$ is costly, one could reasonably argue that operational conservatism may characterize subjects with a poor understanding of the market mechanism. Column 4 of Table 5 includes the number of mistakes made by the subjects in the quizzes and a dummy capturing their (self-reported) low degree of financial literacy. The coefficients of interest are robust to the inclusion of the additional controls.

The consistency between the parameters corroborates the conjecture that δ_i , despite being inconsistent with expected utility, represents an additional way to express risk preferences. This result is intriguing in our opinion, as it may help explaining why elicited risk attitudes usually have a very limited predictive power (see Isaac and James (2000), Deck, Lee, Reyes, and Rosen (2013), Friedman, Isaac, James, and Sunder (2014),

³²The Spearman correlation is based on the rank order of the coefficients and does not depend on the CRRA assumption. An alternative is to compare θ_i^{mkt} with the rough choice in the Investment Game rather than with θ_i^{inv} , which is the route taken in the regressions that follow.

TABLE 5. Consistency of measures of risk aversion.

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

Note: The table reports regressions on the relation between the choice in the Investment game (dep. variable) and the estimated CRRA coefficient θ_i^{mkt} and operational conservatism coefficient δ_i . “Errors” represent the number of mistakes made before solving the control quizzes. “Low financial literacy” is a dummy taking value 1 if the subject answered to the question “What best describes your knowledge about financial markets?” selecting “poor knowledge.” In parentheses, we report robust standard errors, clustered at the market level.

among others). While economists restrict the risk-aversion concept to the diminishing marginal utility of wealth, subjects are likely to hold a broader representation of this construct.³³

The interpretation of operational conservatism as an expression of risk aversion also allows us to reconcile the individual characteristics of the traders with the market outcomes. We have seen before (Result 2) that market prices do not differ between high and low risk-aversion markets. Consistently with the pattern emphasized in Table 5, subjects in high risk-aversion markets display a higher δ_i (median: 0.70 vs. 0.58, Mann-Whitney U test: $p = 0.061$) and a higher θ_i^{mkt} (median 1.95 vs. 1.72, Mann-Whitney U test: $p = 0.078$). While both δ_i and θ_i^{mkt} imply a lower exposure in the market, they have opposite effects on prices. A higher θ_i^{mkt} tends to push prices away from 50, while a higher δ_i counterbalances this effect. As a consequence, despite the marked differences in individual behavior driving Result 1, aggregate outcomes are similar in low and high risk-aversion markets (Results 2 and 3). The positive correlation between θ_i^{mkt} and δ_i across high and low risk-aversion markets emerges despite the two parameters are substitutes at the individual level (Spearman’s $\rho = -0.27$, $p < 0.001$). Subjects in high risk-aversion markets have, by construction, a higher θ_i^{inv} , which they express through either a higher θ_i^{mkt} or through a higher δ_i .

6. CONCLUSION

The role played by risk aversion in information aggregation has been insufficiently investigated empirically. In this paper, we try to fill this gap analyzing experimentally how the prices in a common-value call market respond to risk aversion.

We find that at first glance prices do not react to the level risk aversion, while they should do so according to the theoretical predictions. This result is not due to subjects

³³Holzmeister and Stefan (2021) reach a similar conclusion showing that subjects are aware that the riskiness of their decisions varies across risk elicitation methods.

with different (elicited) risk preferences behaving in a similar manner. We identify two behavioral traits that correlate with elicited risk aversion. The first follows from the traditional curvature of the utility function, which maps into the slope of the demand schedule. The second is the inclination to act as if one possessed less information than he does, which maps into the intercept of the demand schedule. We label this tendency *operational conservatism*. These two facets of cautious behavior have countervailing effects on prices, explaining our aggregate results.

Information aggregation is typically tested using double auctions since call markets are known to be an ineffective architecture (Kagel and Levin (1986), Guarnaschelli, Kwasnica, and Plott (2003), Chen and Plott (2008)). We show, indeed, that the information brought by traders in the market is only part of the information possessed. While Page and Siemroth (2021) come to similar conclusions for double-auction experiments, by reverse-engineering the observed prices, we do so by looking directly at demand schedules. These determine prices that are observationally similar to those that would have been observed under risk neutrality and aggregation of full information. A natural development of our paper is then to extend the analysis to the double auction, and investigate the role of risk preferences within an institution that is relatively more effective and where risk preferences may interact with inferences on others' private information derived from observed prices.

On the side of the measurement of risk preferences, our results do not add to the long list of contributions claiming the empirical failure of the assessment of risk attitudes when not of expected utility theory as a whole. By hypothesizing that *operational conservatism* is part of subjects' representation of risk aversion, our evidence is more optimistic about the possibility to capture stable features of choice under risk even across different contexts and including a complex environment like a call market. In order to achieve such a goal, it is necessary to embrace a broader representation of risk attitudes than implied by diminishing marginal utility of wealth. A natural extension of our study is to analyze whether *operational conservatism* extends to other environments. The evidence of our experiment is promising, but further tests are needed in order to fully identify its nature, implications, and relation with other individual characteristics.

APPENDIX A: EXPERIMENTAL INSTRUCTIONS

Welcome to this experiment and thank you very much for taking the time to support our research. In the next 2 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 $100 \text{ MU} = 1 \text{ euro}$.

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. Noninvested 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 a Head shows up, the investment is successful and you receive the amount not invested plus 2.5 times the investment;
- If a 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.

[Only in CA: 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 2 minutes each.]

Let us now answer in detail 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

TABLE A.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	-	-	-	-	-
$p(\text{Blue} s)$	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

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 in Table A.1.

You receive with the same probability one of the possible signals given the selected urn.

[CA: Signals are randomly assigned and are different for each participant in the same market.]

[RPM: Signals are randomly assigned so that all the 11 signals given the urn are distributed to 2 participants in this session.]

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.

[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 held 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 toward 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?

[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).]

[RPM: In every period, the market stays open for 2 minutes after which the trading system computes the market price for every participant as explained 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]

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.

TABLE A.2. Example of trading activity and final earnings

	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

[CA: Consider the example in Table A.2: 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:]

[RPM: Consider the example in Table A.2: Let us evaluate the situation of a buyer who buys 10 assets and of seller selling 10 assets at a market price of 50 MU:]

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, that is, 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 four 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).

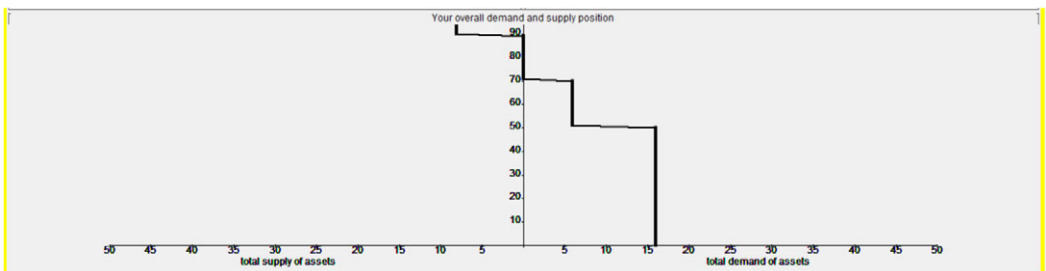


(c) "Book": it shows all your orders to buy (left) and to sell (right) the asset.

Your buy orders		Your sell orders	
Quantity	Price	Quantity	Price
6	70	8	90
10	50		

[DELETE](#) [DELETE](#)

(d) "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

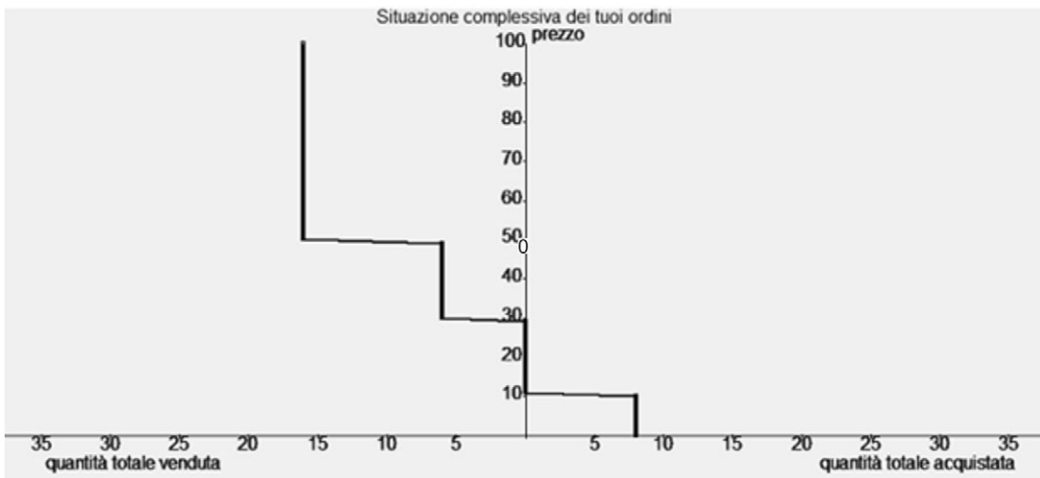
Each order has a limit price but is also executed for “better” prices.

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, that is, 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 us make another example with the following orders:

Tui ordini in acquisto		Tui ordini in vendita	
Numero titoli	Prezzo unitario	Numero titoli	Prezzo unitario
8	10	6	30
		10	50
<input type="button" value="ELIMINA"/>		<input type="button" value="ELIMINA"/>	

In this case, you buy 8 units at any price lower than or equal to 10, while you do not 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, that is, 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, and 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 have done all three of them, press CONFIRM.

[QUIZ 5]

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]


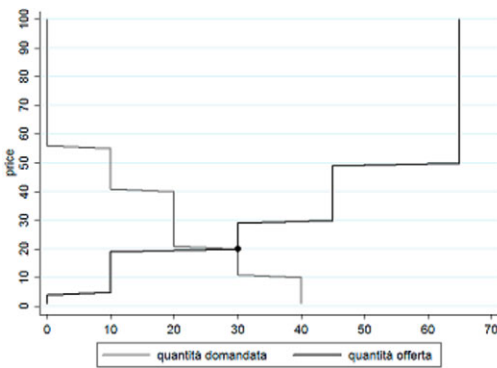
How the market price is determined [CA]

- (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.

How the market price is determined [RPM]

At the end of the 2 minutes for inserting the orders, the system sets a price for each participant. The selected price does NOT depend on the orders entered by the participants, but is randomly drawn according to the following rules. First, the selects a price interval. Different intervals have a different probability of being selected, as shown in Table A.3.

Once the interval has been identified, the system randomly selects one of the prices in that interval (i.e., every price within each interval has the same probability of being selected).

TABLE A.3. Probability distribution of the different price intervals

Price window	1–14	15–29	30–44	45–55	56–70	71–85	86–99
Probability	0.01	0.09	0.20	0.40	0.20	0.09	0.01

Order execution

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.]

Outcome of the trading period

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

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 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.

APPENDIX B: CARA UTILITY FUNCTION

Assuming that traders' preferences are represented by an individual utility function over wealth levels featuring *constant absolute risk aversion* (γ_i):

$$u_i(w_i) = -e^{-w_i\gamma_i} \quad \text{if } \gamma_i > 0$$

it can be shown that under PIE the optimal individual net demand is equal to

$$q_i^*(p, b_i, \gamma_i > 0) = \frac{1}{\gamma_i} \left[\ln\left(\frac{b_i}{p}\right) - \ln\left(\frac{1-b_i}{1-p}\right) \right]. \tag{5}$$

Equation (5) implies that, as in the case of the CRRA (equation (2)):

1. $q_i^* = 0$ if $b_i = p$, that is, the individual demand pivots at the individual belief;
2. $q_i^* > 0$ if $b_i > p$, that is, the trader is a net buyer when the belief is higher than the price, and viceversa;
3. $\delta q_i^* / \delta \gamma_i < 0$, that is, the amount exchanged for any given price and belief decreases with the degree of risk aversion (Hypothesis 1).

The main difference as compared to the CRRA case is that with the CARA the market price under PIE does not deliver the average belief at $\gamma = 1$. In our setting, this condition is ensured by an absolute level of risk aversion $\gamma \cong 0.1$.

Hypotheses 2 and 3 are described in Figure B.1, which replicates the CRRA case illustrated in Figure 2.

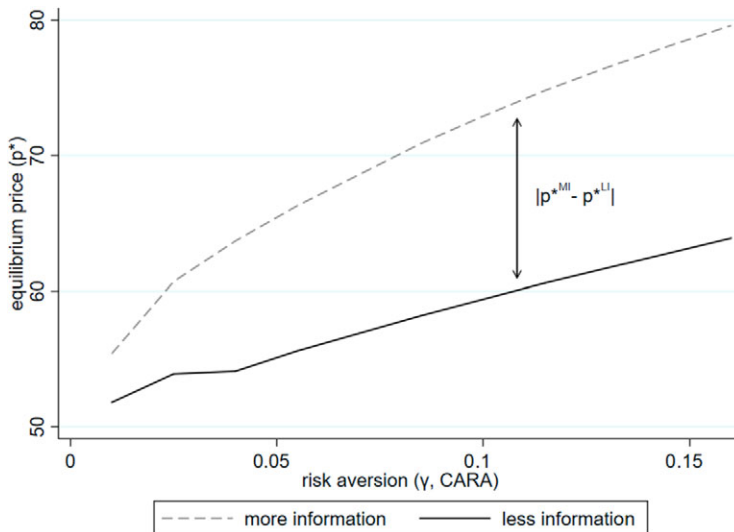


FIGURE B.1. Visual representation of Hypotheses 2 and 3, CARA. *Note:* The figure shows the equilibrium prices for increasing degrees of risk aversion in markets characterized by different levels of information (MI vs. LI). The prices are derived computationally using the empirical distributions of beliefs implemented in our experiment. The distance of each of the two curves from 50 illustrates Hypothesis 2. The distance between MI and LI describes Hypothesis 3.

APPENDIX C: FURTHER RESULTS

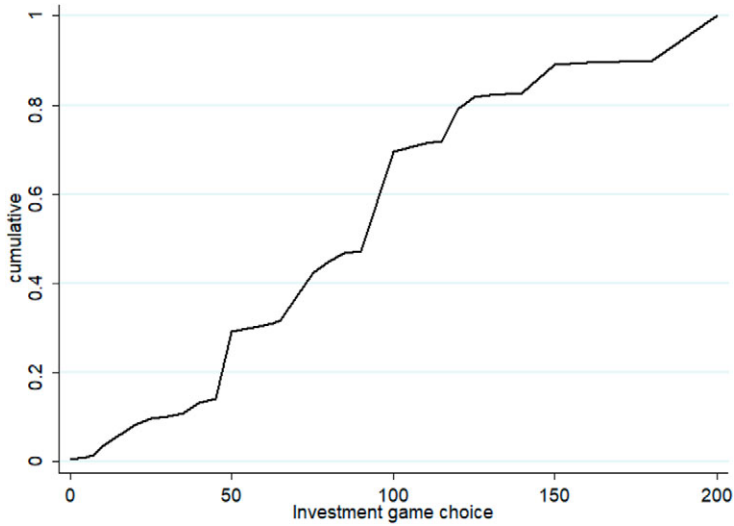


FIGURE C.1. Cumulative distribution of choices in the risk elicitation task.

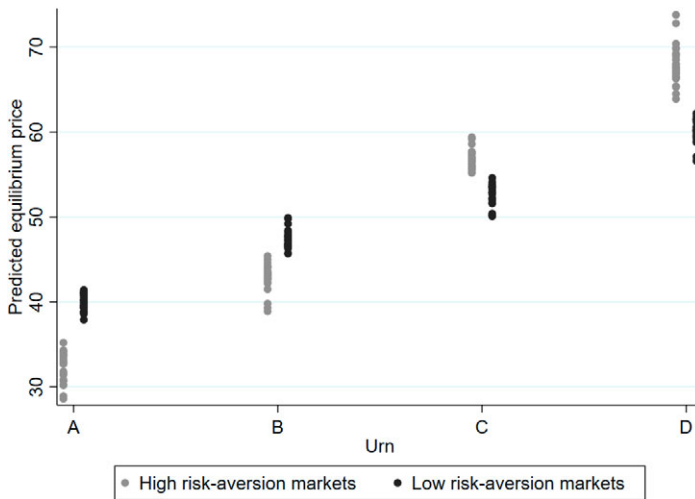


FIGURE C.2. Predicted equilibrium prices, given the joint distribution of signals and CRRA. *Notes:* The figure shows for each urn the average predicted equilibrium price in each market, distinguishing between low and high risk-aversion markets. Predicted equilibrium prices are calculated considering the realized joint distribution of signals (i.e., Bayesian beliefs) and CRRA coefficients as elicited from the Investment Game (θ_t^{inv}), and assuming optimizing behavior according to the prior information equilibrium.




FIGURE C.3. Distance of price from average (elicited) belief and risk aversion. *Notes:* The figure plots the distance of market prices in each market/period from the average elicited belief, plotted against the average CRRA coefficient in the market (the gap between 0.27 and 0.58 is the effect of the matching procedure). A linear fit between the two measures is overimposed.

REFERENCES

- Alti, Aydoğ̃an and Paul C. Tetlock (2014), “Biased beliefs, asset prices, and investment: A structural approach.” *The Journal of Finance*, 69, 325–361. [753]
- Ang, James S. and Thomas Schwarz (1985), “Risk a version and information structure: An experimental study of price variability in the securities markets.” *The Journal of Finance*, 40, 825–844. [755]
- Arrow, Kenneth J., Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson et al. (2008), “The promise of prediction markets.” *Science*, 320, 877–878. [753]
- Asparouhova, Elena, Peter Bossaerts, Jon Eguia, and William Zame (2015), “Asset pricing and asymmetric reasoning.” *Journal of Political Economy*, 123, 66–122. [775]
- Becker, Gordon M., Morris DeGroot, and Jacob Marschak (1964), “Measuring utility by a single-response sequential method.” *Behavioral Science*, 9, 226–236. [766]
- Biais, Bruno, Thomas Mariotti, Sophie Moinas, and Sébastien Pouget (2017), *Asset Pricing and Risk Sharing in a Complete Market: An Experimental Investigation*. Technical report, Toulouse School of Economics (TSE). [755]
- Bliss, Robert R. and Nikolaos Panigirtzoglou (2004), “Option-implied risk aversion estimates.” *The journal of finance*, 59, 407–446. [756]
- Bossaerts, Peter, Cary Frydman, and John Ledyard (2013), “The speed of information revelation and eventual price quality in markets with insiders: Comparing two theories.” *Review of Finance*, 18, 1–22. [765]

Bossaerts, Peter and William R. Zame (2008), "Risk aversion in laboratory asset markets." In *Risk Aversion in Experiments*, 341–358, Emerald Group Publishing Limited, Bingley. [755]

Chen, Kay-Yut and Charles R. Plott (2008), "Markets and information aggregation mechanisms." *Handbook of Experimental Economics Results*, 1, 344–352. [781]

Choo, Lawrence, Todd R. Kaplan, and Ro'i Zultan (2017), "Information aggregation in arrow–Debreu markets: An experiment." *Experimental Economics*, 1–28. [772]

Cipriani, Marco and Antonio Guarino (2014), "Estimating a structural model of herd behavior in financial markets." *American Economic Review*, 104, 224–251. [753]

Cox, James C., Bruce Roberson, and Vernon L. Smith (1982), "Theory and behavior of single object auctions." In *Research in Experimental Economics*, Vol. 2 (Vernon L. Smith, ed.), 1–43, JAI Press, Greenwich. [756]

Crockett, Sean, Daniel Friedman, and Ryan Oprea (2021), "Naturally occurring preferences and general equilibrium: A laboratory study." *International Economic Review*, 62, 831–859. [754]

Crosetto, Paolo and Antonio Filippin (2013), "The 'bomb' risk elicitation task." *Journal of Risk and Uncertainty*, 47, 31–65. [763]

Crosetto, Paolo and Antonio Filippin (2016), "A theoretical and experimental appraisal of four risk elicitation methods." *Experimental Economics*, 19, 613–641. [762, 769]

Deck, Cary, Jungmin Lee, Javier A. Reyes, and Christopher C. Rosen (2013), "A failed attempt to explain within subject variation in risk taking behavior using domain specific risk attitudes." *Journal of Economic Behavior & Organization*, 87, 1–24. [779]

Dimmock, Stephen G., Roy Kouwenberg, and Peter P. Wakker (2016), "Ambiguity attitudes in a large representative sample." *Management Science*, 62, 1363–1380. [775]

Easley, David, Nicholas M. Kiefer, and Maureen O'Hara (1997), "One day in the life of a very common stock." *The Review of Financial Studies*, 10, 805–835. <http://dx.doi.org/10.1093/rfs/10.3.805>. [753]

Enke, Benjamin and Thomas Graeber (2019), *Cognitive Uncertainty*. Technical report, National Bureau of Economic Research. [775]

Epstein, Larry G. (2006), "An axiomatic model of non-Bayesian updating." *The Review of Economic Studies*, 73, 413–436. [777]

Fattinger, Felix (2018), "Trading complex risks." Available at SSRN 3086358. [774, 775]

Fellner, Gerlinde and Boris Maciejovsky (2007), "Risk attitude and market behavior: Evidence from experimental asset markets." *Journal of Economic Psychology*, 28, 338–350. [755]

Filippin, Antonio and Marco Mantovani (2023), "Supplement to 'Risk aversion and information aggregation in binary-asset markets.'" *Quantitative Economics Supplemental Material*, 14, <https://doi.org/10.3982/QE1981>. [753]

- Fischbacher, Urs (2007), “z-tree: Zurich toolbox for ready-made economic experiments.” *Experimental Economics*, 10, 171–178. [766]
- Fountain, John and Glenn Harrison (2011), “What do prediction markets predict?” *Applied Economics Letters*, 18, 267–272. [758]
- Friedman, Daniel, R. Mark Isaac, Duncan James, and Shyam Sunder (2014), “Risky curves: On the empirical failure of expected utility.” Routledge. [779]
- Fryer, Roland G. Jr. (2013), *Information and Student Achievement: Evidence From a Cellular Phone Experiment*. Technical report. National Bureau of Economic Research. [776]
- Gjerstad, Steven (2005), “Risk aversion, beliefs, and prediction market equilibrium.” *Microeconomics* 0411002, EconWPA. [754, 758, 759]
- Gneezy, Uri and Jan Potters (1997), “An experiment on risk taking and evaluation periods.” *The Quarterly Journal of Economics*, 112, 631–645. [754, 762]
- Guarnaschelli, Serena, Anthony M. Kwasnica, and Charles R. Plott (2003), “Information aggregation in double auctions: Rational expectations and the winner’s curse.” *Information Systems Frontiers*, 5, 63–77. [781]
- He, Xue-Zhong and Nicolas Treich (2017), “Prediction market prices under risk aversion and heterogeneous beliefs.” *Journal of Mathematical Economics*, 70, 105–114. [758, 761]
- Hellwig, Martin F. (1980), “On the aggregation of information in competitive markets.” *Journal of economic theory*, 22, 477–498. [761]
- Holt, Charles A. and Susan K. Laury (2002), “Risk aversion and incentive effects.” *American Economic Review*, 92, 1644–1655. [763]
- Holzmeister, Felix and Matthias Stefan (2021), “The risk elicitation puzzle revisited: Across-methods (in) consistency?” *Experimental Economics*, 24 (2), 593–616. [780]
- Horn, Christian Franz and Björn Sven Ivens (2015), “Corporate prediction markets for innovation management.” In *Adoption of Innovation*, 11–23, Springer. [754]
- Hossain, Tanjim and Ryo Okui (2013), “The binarized scoring rule.” *The Review of Economic Studies*, 80, 984–1001. [764]
- Isaac, R. Mark and Duncan James (2000), “Just who are you calling risk averse?” *Journal of Risk and Uncertainty*, 20, 177–187. [779]
- Kagel, John H. and Dan Levin (1986), “The winner’s curse and public information in common value auctions.” *The American economic review*, 76 (5), 894–920. [781]
- Luckner, Stefan, Jan Schröder, and Christian Slamka (2008), “On the forecast accuracy of sports prediction markets.” In *Negotiation, Auctions, and Market Engineering*, 227–234, Springer. [754]
- Manski, Charles F. (2006), “Interpreting the predictions of prediction markets.” *Economics Letters*, 91, 425–429. [754, 759]

Milgrom, Paul and Nancy Stokey (1982), "Information, trade and common knowledge." *Journal of Economic Theory*, 26, 17–27. [754]

Ngangoué, M. Kathleen and Georg Weizsäcker (2021), "Learning from unrealized versus realized prices." *American Economic Journal: Microeconomics*, 13, 174–201. [755]

Page, Lionel and Christoph Siemroth (2021), "How much information is incorporated into financial asset prices? Experimental evidence." *The Review of Financial Studies*, 34, 4412–4449. [755, 776, 781]

Pennings, Joost M. E. and Ale Smidts (2000), "Assessing the construct validity of risk attitude." *Management Science*, 46, 1337–1348. [755]

Peterson, Cameron R. and Alan J. Miller (1965), "Sensitivity of subjective probability revision." *Journal of Experimental Psychology*, 70, 117–121. [774]

Plott, Charles R. and Shyam Sunder (1982), "Efficiency of experimental security markets with insider information: An application of rational-expectations models." *Journal of political economy*, 90, 663–698. [772]

Plott, Charles R. and Shyam Sunder (1988), "Rational expectations and the aggregation of diverse information in laboratory security markets." *Econometrica: Journal of the Econometric Society*, 56 (5), 1085–1118. [772]

Snowberg, Erik and Justin Wolfers (2010), "Explaining the favorite–long shot bias: Is it risk-love or misperceptions?" *Journal of Political Economy*, 118, 723–746. [759]

Trautmann, Stefan T. and Gijs Van De Kuilen (2015), "Ambiguity attitudes." *The Wiley Blackwell handbook of judgment and decision making*, 1, 89–116. [775]

Wolfers, Justin and Eric Zitzewitz (2004), "Prediction markets." *Journal of economic perspectives*, 18, 107–126. [753]

Wolfers, Justin and Eric Zitzewitz (2006), *Interpreting Prediction Market Prices as Probabilities*. Technical report. National Bureau of Economic Research. [759]

Co-editor Garance Genicot handled this manuscript.

Manuscript received 13 September, 2021; final version accepted 5 October, 2022; available online 11 October, 2022.