Communication islands, biased information sources, and asset markets*

Anna Bayona[†], Jordi Brandts[‡] and Xavier Vives[§]

May 27, 2026

Abstract

We design a laboratory asset-market experiment to examine how the aggregation of traders' private information about the asset's value depends on two factors: (i) the bias in their information sources and (ii) the structure of the communication network. If there is bias, half of the participants receive their signals from one source and the other half from another source, with the two sources being biased in different directions. Participants interact under one of three communication regimes: full communication, two communication groups segregated by signal source, or no communication. We find that communication leads to prices being more sensitive to fundamental values than its absence, regardless of the bias. When information sources are biased, we find that segregated communication produces more informative prices than full communication. Belief elicitation mirrors these price patterns. Analysis of chat content suggests a mechanism: segregation increases the frequency of true consensus about the asset's fundamental value, improving belief accuracy. Full communication, by contrast, aggregates effectively when the information structure is unbiased. However, under bias, it mixes two oppositely biased sources into a conversation traders find harder to interpret.

JEL: C92, D82, D83, D84, G10, G14
Keywords: information, market efficiency, networks, experimental asset market.

*We thank participants at 2025 IMEBESS (Valencia), the 2025 Experimental Finance Conference (Maastricht), the 2025 SAET Conference (Ischia), and the 2026 IAREP Conference (Castellón de la Plana). We also thank Penélope Hernández, Aurora García-Gallego, Manuel Genovés and Manuel Guerrero for excellent support in conducting the experimental sessions, and Daniel Carrera, Apurva Ashish Tudekar and Martín Villarán for excellent research assistance. The authors gratefully acknowledge financial support from the Fundació d'Economia Analítica under the research agreement with IESE Business School. Bayona acknowledges the financial support from the Spanish Ministry of Science, Innovation and Universities and Agencia Nacional de Investigación (Grant: PID2024-156339NA-I00) financed by FEDER and EU. Brandts thanks the Spanish Ministry of Economy and Innovation (Grant: PID2023-147136NB-I00) and the Severo Ochoa Program for Centers of Excellence in R&D (CEX2019-000915-S).

[†]ESADE Business School, Universitat Ramon Llull. E-mail: anna.bayona@esade.edu.
[‡]IAE-CSIC and Barcelona School of Economics. E-mail: jordi.brandts@iae.csic.es.
[§]IESE Business School. E-mail: xvives@iese.edu.

1 Introduction

Online communication has become an important source of information in a variety of economic, social and political settings. It has drastically changed the way individuals share relevant (and irrelevant) information with each other and is susceptible of affecting their beliefs and decisions.

In this paper we present the results from a series of laboratory experiments related to how information transmission through communication networks affects asset markets. Information transmission through such channels may lead to quicker dissemination of private information to the market. However, it may also worsen information transmission. Hirshleifer (2020) argues that most of the effects of social media in finance occur due to the "social transmission bias" which can take several forms. Financial information that flows through electronic channels is prone to be influenced by the facts that different traders obtain their information about the market from different sources and communicate with other traders on separate communication islands. In this paper we study whether information transmission under such conditions affects behavior in a laboratory experiment asset market.

We examine an environment with a single asset, two states of the world where the value of the fundamental migh be high or low. Participants in the experiments do not know the state of nature before trading, and each participant receives a private signal about it. After receiving the private signal, participants may communicate in a free-form chat with others through their connections in a communication network. Then, we elicit participants' beliefs about whether the state is high or low. After this stage, all agents trade in a continuous double auction, which resembles trading in real financial markets. Specifically, our study focuses on how financial market outcomes are jointly affected by: (i) the variation of communication network structures; (ii) the variation of the information structures from which traders obtain their private signals. For this purpose, we conduct a 2×3 design.

The first dimension we vary is the information structure from which traders obtain their private signals. An information structure specifies, for each state of the world, a probability distribution over the signals a participant may receive. We consider two cases. In the unbiased case, all traders draw signals from the same, unbiased information structure, which is equally informative in both states. Following Gentzkow *et al.* (2015) and Charness *et al.* (2021), in the biased case, there are two information structures that are biased in opposite directions: one is biased toward one state and the other toward the opposite state. Half of the traders receive their signals from one structure and half from the other, so that the two halves are informed by oppositely biased sources. The two biased structures are symmetric — biased to the same degree but in opposite directions — and are parametrized so that the unbiased structure is as informative about the true state as the average of the two biased ones. This holds the overall informativeness of the environment fixed across the two cases, isolating the effect of distortion itself.

The second dimension we vary is the communication network structure. We consider three cases: a baseline in which no communication between traders is possible; a second network in which all traders can communicate with one another before trading; and a third in which traders are divided into two equal, disjoint communication islands. In the treatment with biased signals and two disjoint islands, all members of one island receive signals from the structure biased toward one state, while all members of the other island receive signals from the structure biased toward the opposite state. In the two cases with communication, all traders have identical communication opportunities in a free-form chat. These structures were chosen for their simplicity and for the sharp contrast they draw between full connectedness and segregation into communication islands.

With our design we capture the simple notions that some traders are sometimes, due to biased signals, better or worse informed than others and that the better informed and the worse informed may communicate all together or in separate groups. The above variations give rise to a 3x2 design which allows us to study the interaction between communication network and information structures.

Communication between traders involves a trade-off with respect to its effect on market performance. On the one hand, it allows private information to disseminate beyond what trading itself reveals through bids, asks, and prices, thereby improving market conditions. On the other hand, it can also give rise to the social transmission bias proposed by Hirshleifer (2020). In our setting, this bias has two distinct sources. First, it can arise from the signals being biased at the source, independently of how participants communicate. Second, it can arise in the communication stage itself, since in free-form chat participants may lie about their signal or, more broadly, use natural language in misleading ways. Our design lets us separate these two channels and study their interaction: we can examine how biased signal structures and participants' scope for misrepresentation jointly shape market outcomes across different communication networks. The experimental approach makes it possible to identify precisely how communication and information structures jointly affect market behavior.

Our fundamental research questions are the following: Does communication lead to different market behavior relative to the baseline of no communication? Is there a difference in the above between communication taking place on only one island or on two disjoint islands? Is there a difference in the above between signals being biased and unbiased? Are there interaction effects between communication and information conditions?

The main focus of our analysis is on the deviations of prices from the fundamental value of the asset, along with other measures of market quality, although we will also examine how other aspects of our results — such as traders' beliefs and communication content — are affected by communication networks and information structures. Because we use a laboratory experiment, we have control over the fundamental value of the asset and can therefore observe directly the difference between transaction prices and the fundamental value. At the same time, we can study other features of market behavior within a unified framework.

We report four main findings. First, prices track fundamentals only imperfectly in every treatment: low-value assets trade above their worth and high-value assets below it. Communication helps — markets with chat respond far more strongly to fundamentals than silent ones — but it never delivers full efficiency.

Second, the value of communication depends on the type of the underlying signals, and the two interact. When signals are biased, segregating traders into separate islands yields the most accurate prices and beliefs; when signals are unbiased, connecting everyone works best. Full communication with biased signals is the worst case of all, leaving beliefs no more accurate than if traders had not communicated at all. We also show that the main results do not depend on whether we display the posterior probabilities to participants.

Third, this interaction operates through what traders come to believe. Segregation improves market accuracy mainly by helping groups converge on correct shared views: relative to full communication, segregated chat produces more frequent true agreement without producing any more false agreement. The structure of the network — not merely the presence of communication — is what matters once signals are biased.

Fourth, liquidity responds to network structure differently depending on the information environment, and the two liquidity measures move in different conditions: full communication dampens trading volume when signals are unbiased, whereas segregation widens bid–ask spreads when signals are biased. Trading intensity and trading cost are thus distinct dimensions rather than facets of a single notion of liquidity.

Our work relates to the broad area of social finance that uses field data. Kuchler and Stroebel (2021) review empirical literature based on field data that studies how social interactions affects financial decisions of households, retail investors and professional investors. Cookson *et al.* (2024) present a broad overview of research on social media and finance in which they cover contributions related to the effects on the informational efficiency of prices and argue that whether social media helps or hurts market efficiency is an open question. Keasey *et al.* (2025) find that mega influencers affect investors' attention, volatility and trading volume but not stock returns. Only posts by top influencers with extreme sentiment affect returns and, even here, the effect is short-lived. Chen *et al.* (2014) analyze the content of articles published on one of the most popular social media platforms for investors in the United

States and find that the views expressed in both articles and commentaries predict future stock returns and earnings surprises. Jiao *et al.* (2020) find that coverage by traditional news media predicts decreases in subsequent volatility and turnover, but coverage by social media predicts increases in volatility and turnover and show that these patters are consistent with a model of "echo chambers", where social networks repeat news, but some investors interpret repeated signals as genuinely new information. Hirshleifer *et al.* (2024) study how the social transmission of public news influences investors' beliefs and the securities markets and find that earnings announcements from firms in higher-centrality counties generate a stronger immediate price, volatility, and trading volume reactions. Their findings suggest that greater social connectedness promotes quick incorporation of news into prices, but also opinion divergence and excessive trading.

Several studies have tried to theoretically model the role of social communication in financial markets.[1] Choi *et al.* (2017) study a rational expectations model where agents are connected through a network. They find that when information is exogenous, social communication improves market efficiency, however social communication crowds out information acquisition. Walden (2019) model a dynamic rational expectations model where information diffuses through a network deriving implications for trading volume and similarity of trades. Pedersen (2022) builds a theoretical model with four types of investors (naive, fanatic, rational short term and rational long term) that coexist in a network, share information and then trade.[2]

There is a long history of using laboratory experiments to study issues of information aggregation and dissemination in experimental asset markets (Forsythe *et al.*, 1982; Friedman

[1] Acemoglu *et al.* (2024) study a model of online content sharing, where articles may or not contain misinformation. They find that homophily reduces the circulation of an article but it also creates echo chambers that have less discipline on sharing content which is less reliable.

[2] Goldstein *et al.* (2025) study information sharing in a Kyle model where investors can be informed or uninformed. They show that the uninformed investor has incentives to share information with the informed, which leads the informed investor to trade against the uninformed investors' information, thereby reducing his price impact. Paradoxically, the well informed investor loses from receiving information because of the resulting worsened market liquidity and the more aggressive trading by the coarsely informed investor. However, this model does not incorporate a network or signal distortions.

et al., 1984; Plott and Sunder, 1982; Sunder *et al.*, 1992; Plott and Sunder, 1988). Our paper focuses on how the market aggregates information by diverse traders and how this information is disseminated to all traders. The first seminal paper which focuses on this is Plott and Sunder (1988). They find that aggregation of information depends on market features, with results ranging from no aggregation to full aggregation. Corgnet *et al.* (2022) replicate the results of Plott and Sunder (1988) and find that information aggregation in asset markets is fragile and only occurs in very special environments.

A few experimental studies focus on traders' information acquisition in experimental markets. Halim *et al.* (2019) test experimentally the predictions of Han and Yang (2013). They find that social communication crowds out information acquisition since agents may free ride on the information of their neighbors, and consequently, social communication increases trading volume and liquidity but does not improve price informativeness. In these models, social communication is non-strategic. Page and Siemroth (2017) study individual information acquisition in experimental prediction markets. They find overacquisition of information leading to traders obtaining negative profits net of information costs and propose this finding as a novel explanation for the high forecasting accuracy of prediction markets. Page and Siemroth (2021) find that private information dispersed across traders is substantially underincorporated (less than 50%) into prices of prediction markets. Corgnet *et al.* (2024) present the results from experimental asset markets in which communication between traders improves informational efficiency. They do not study the effects of segregated communication islands and signal distortion, which are the main focus of the work we present in this paper. In addition, in all our treatments communication is free-form and we study the relation between communication content and market outcomes.

This paper is organized as follows. In Section 2, we present the experimental design. In Section 3, we analyze the main results of the experiment. Section 4 we discuss potential mechanisms. Section 4 provides a discussion; and Section 5 concludes. The Appendix provides full instructions for one treatment, details of the chat analysis, and further analysis of our results. The Online Appendix provides full experimental instructions on all treatments.

2 Experimental Design

We conduct a pre-registered experiment with six treatments that vary along two dimensions in a 2×3 between-subject design.[3]

The first dimension concerns the information source from which participants receive private signals about the asset value, which has two levels: (1) unbiased signals (U); and (2) biased signals (B). For this part of the design, we build on Charness *et al.* (2021), particularly for the biased-information case. In the experiment, we present these sources as "advisors" who transmit information to participants. In the non-biased case, all participants receive signals from the same advisor, who draws signals for different traders from a common probability distribution. In the biased case, participants are assigned to one of two advisors: the signals of one advisor are biased in one direction, while those of the other are biased in the opposite direction. This representation of biased signals can drive a wedge between the beliefs of traders who receive information from different advisors.

The second dimension concerns the communication network, which has three levels: (a) no communication (NoCom); (b) segmented communication between two disjoint groups of traders (SegCom); and (c) full communication among all traders (AllCom). This yields the six treatments summarized in Table 1, together with their acronyms.

Table 1
Treatment summary

Communication network	Unbiased Signals	Biased signals
No communication	U-NoCom	B-NoCom
Segregated communication	U-Seg	B-Seg
All communication	U-AllCom	B-AllCom

The structure of the experimental game that participants play in our experiment is as follows. Each market consists of 8 participants that are endowed with cash and several units of a risky asset of uncertain common value (the fundamental or state). Each participant

[3] The experiment received the Ethics Approval of the ESADE Ethics Committee (019/2024), and the experiment was pre-registered at AsPredicted with registration number 196179 and pdf: https://aspredicted.org/cv5kp3.pdf

receives a private signal about the asset value. After receiving the private signal, participants may communicate (in the communication treatments) through a given network that allows free form communication among neighbours. We then elicit participants' beliefs about the fundamental value of the asset. Then, agents can trade in a continuous double auction for several minutes. Payoffs are realized. Traders' private beliefs about the fundamental value of the asset are elicited again. Participants receive feedback and then the next round starts. They play this game for 12 rounds. There is feedback between rounds. After these rounds were completed, participants answered several post-experimental questionnaires. Before playing the main rounds, participants completed 2 trial rounds to foster understanding of the environment.

We now proceed to explain the details of each of the three parts of the experiment.[4]

Part I: Information, (Chat), and Beliefs

In Part I, participants receive a binary private signal about the financial asset; in the four communication treatments, they can also communicate with other participants via free chat; afterwards, we elicit their probabilistic beliefs about whether the asset value is high or low.

The asset's value is the same for all participants in a market. It is determined by the computer drawing either an orange (O) or a green (G) ball from an urn. If the ball is orange, the asset is worth 50 ECUs in that round; if it is green, it is worth 500 ECUs. At the start of each round, participants are told the prior probability of the color of the ball. It is explained to subjects that balls are drawn from one of the two possible urns: urn 1, containing 6 orange and 4 green balls, or urn 2, containing 4 orange and 6 green balls. Participants do not observe the ball's color. The urn alternates across rounds, and within each urn every ball is equally likely to be drawn, with replacement. Participants are told that their objective is to infer the ball's color, since this determines the asset's payoff at the end of Part II.

In each round, participants receive a binary private signal about the color of the ball. We

[4] Appendix A contains the complete instructions for Biased Information and Segmented Communication treatments, and the instructions for the rest of the treatments can be found in the Online Appendix.

explain this in terms of an "advisor". Two computerized advisors (A and B) provide private information about the color of the ball to participants. Participants have the same advisor in all rounds. The structure of this private information is a key design feature. In each market, four of the eight participants receive advice from advisor A and the other four from advisor B, with a distribution of advisors to participants that depends on the treatment. Both advisors provide imperfect information that depends on the realized ball color, O or G.

In the unbiased information treatments, the two advisors provide identical, unbiased advice. Thus, all participants effectively receive information from the same source. However, for parallelism with the biased-information case, where the two advisors differ, the instructions still present advisor A and advisor B separately and half of the eight participants receive information from one of the advisors and the other half from the other. In these treatments, the advisors' signals follow the information structure below where θ denotes the state (the ball drawn from the urn) and s the private signal:

Table 2
Unbiased information structure.

	$s = o$	$s = g$
$\theta = O$	0.7	0.3
$\theta = G$	0.3	0.7

Note: Each entry gives the probability of signal s conditional on state θ.

In this unbiased case, the signals are equally informative for both states (Orange and Green)

Under biased information advisor A's information follows the structure on the left below, and advisor B's the one on the right of the table below:

Signals remain informative, the information structure on the right hand side (Advisor B) is biased to Green in relation to the information structure that is depicted on the left-hand side (Advisor A).[5] These two structures are symmetrically biased: both advisors are biased to the same degree, but in opposite directions. These two cases allow a comparison between

[5] One possible interpretation of this asymmetry is reluctance to accept evidence in favor of one or the other state.

TABLE 3
BIASED INFORMATION STRUCTURES.

Advisor A

	$s = o$	$s = g$
$\theta = O$	0.8	0.2
$\theta = G$	0.4	0.6

Advisor B

	$s = o$	$s = g$
$\theta = O$	0.6	0.4
$\theta = G$	0.2	0.8

Note: Each entry gives the probability of signal s conditional on state θ. Advisor A is biased toward O and advisor B toward G; the two structures are symmetrically biased, biased to the same degree in opposite directions.

unbiased (unbiased) signals and biased signals. The signal parametrization was chosen so that the unbiased information structure has the same informativeness as the average of the two biased structures. So, this allows us to compare biased information sources to information structures that are equally informative about the true state of nature.

In both the unbiased and biased cases, the instructions display the information structures of both advisor A and advisor B, so participants can observe both their own advisor's structure and the other advisor's structure. The instructions also emphasize that two participants with the same advisor may nonetheless receive different signals, since signals are independently drawn.

In our baseline design, the instructions present next, for both priors (the two urns), the posteriors conditional on the advisor and the advisor's message. Participants are also told that, in each round, they will see on screen the relevant posterior probability given the selected urn and their advisor's signal.

With an unbiased advisor, posteriors are symmetric around the prior.[6] If the prior favors orange (urn 1), then the conditional probability that the ball is orange given that the advisor has provided an orange signal has a greater posterior probability than if the advisor says green; if the prior favors green (urn 2), then the conditional probability that the ball is green given that the advisor has provided an green signal has a greater posterior probability than if the advisor says orange.

[6]See Appendix A2 for the portion of the Unbiased treatments instructions, which explain the information design, and provides the posterior probabilities.

For biased signals, advisors' bias introduces asymmetries in how priors and signals shape posteriors.[7] For advisor A, if the prior leans toward orange (green), then the conditional probability that the ball is orange (green) given that the advisor said orange (green) has a high (much higher) posterior probability than if the advisor had said green (orange). For advisor B, this is reversed symmetrically.

Providing posteriors is an important design choice: it ensures that observed behavior is not affected by participants' inability to update beliefs correctly (see Halim *et al.*, 2019). However, in two robustness treatments we do not inform participants about the posteriors (see section 5).

The assignment of advisors to participants depended on the network structure. In the unbiased-signal treatments, the two advisors were identical, so the assignment was immaterial. In the biased signal treatments, NoCom and AllCom each had half the participants randomly assigned to advisor A and half to advisor B. In SegCom, by contrast, all members of one island received advisor A and all members of the other received advisor B, so that each island consisted of followers of a single advisor — a structure that captures a form of homophily.

In the communication treatments, participants are connected to a set of other participants, called neighbors. In these treatments, participants can communicate with neighbors via a chat in free form (except that disrespectful language and disclosure of identities are prohibited) for 2 minutes. Under all communication (AllCom), participants are fully connected so that each participant has seven neighbours. Under segregated communication, the group of eight is split into two groups of four, so each participant has three neighbours. These groups of four or eight remain fixed throughout the experiment.

Figure 1 provides a visual summary of the treatment structure.

Finally, we elicit participants' pre-trading probabilistic beliefs about the color of the ball.

[7] See Appendix A1 for the instructions of the biased treatments, which explain the information design and provides the posterior probabilities.

Figure 1: **The six experimental treatments.** The figure displays the two-by-three factorial design that combines three communication network structures (*NoCom*, *SegCom*, *AllCom*) with two information conditions (*Unbiased*, *Biased*). Each node represents one trader, and lines indicate the communication links available within the treatment. In every market, eight traders interact; in the biased information condition, four traders follow advisor A and four follow advisor B, with the four–four split held constant across all three network structures. Under *NoCom*, traders cannot communicate. Under *SegCom*, traders are partitioned into two fully connected islands of four; in the biased condition, each island is composed of followers of a single advisor. Under *AllCom*, the eight traders form a fully connected network. The alternating arrangement of advisor types in the *NoCom* and *AllCom* panels is for visual clarity only; in the experiment, advisor assignment within these treatments was random.

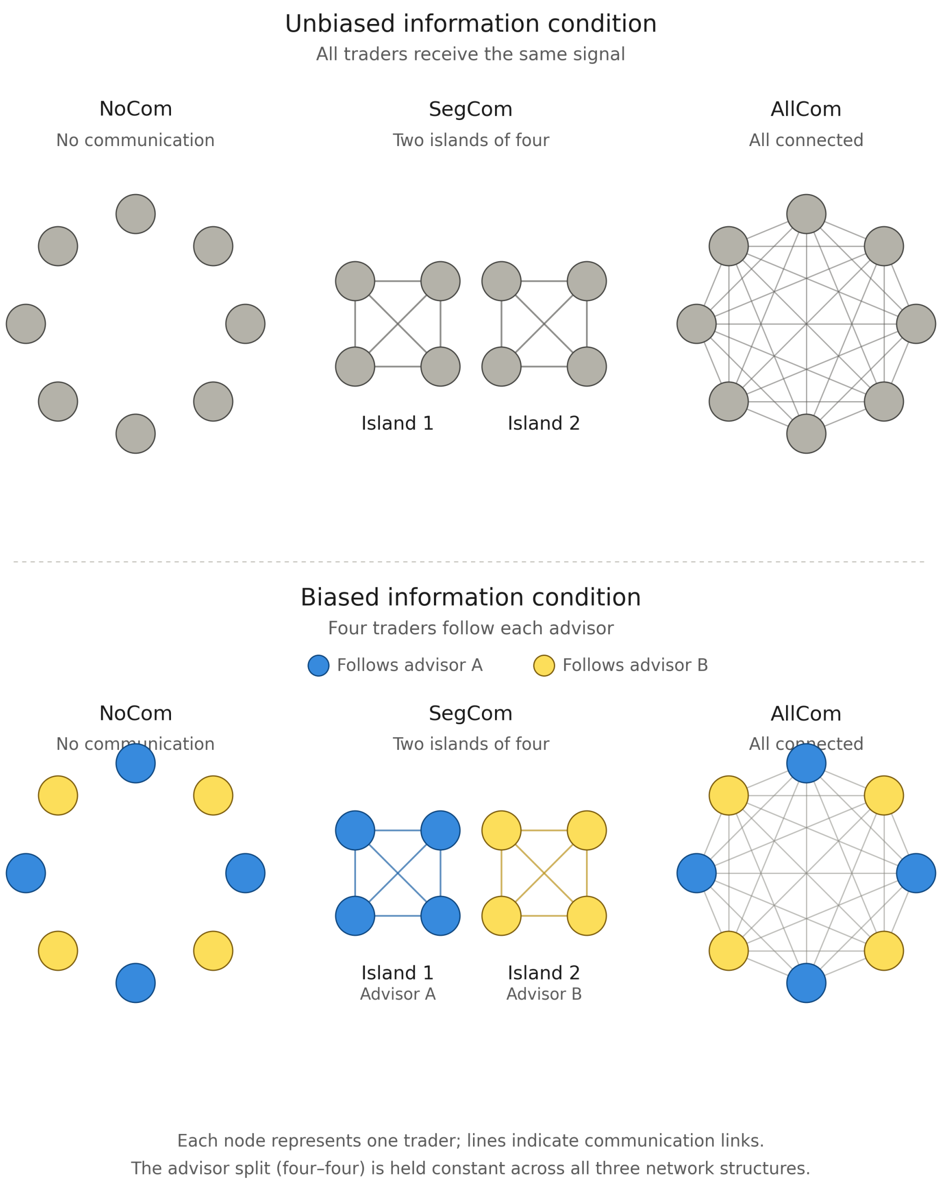

These beliefs are not incentivized.

Part II: Financial Market Trading

In Part II, participants trade in a financial market for 3 minutes per round, with the full group of 8 participants, over 12 independent rounds. In each round, every participant starts with the same endowment: 1000 Experimental Currency Units (ECU) and 4 units of the asset. Round earnings depend on the participant's final ECU holdings, final asset holdings, and the true color of the ball. ECUs and assets do not carry over across rounds.

The market is organized as a continuous double auction, which resembles limit order book trading in many financial markets. To buy one unit of the asset, a participant can either accept a standing ask or submit a bid at a chosen price. To sell one unit, a participant can either accept a standing bid or submit an ask at a chosen price. Buying or selling multiple units requires repeating these actions. All standing bids and asks are visible at all times and are sorted so that bids appear from highest to lowest and asks from lowest to highest, with the best offers shown first. Participants are informed that they should generally accept the highest bids and the lowest asks. Participants cannot accept their own offers, but they can cancel them. Offers disappear once accepted or withdrawn, and also if a participant can no longer execute them because of insufficient assets or ECUs. After each trade, ECU holdings and asset holdings are updated immediately. The remaining time, in seconds, is displayed in the top-right corner of the trading screen. After trading, we elicited participants' post-trading beliefs about the color of the ball.

At the end of the round, a feedback screen reports the average trade price and the participant's profit. The next round begins once all traders click "Next."

Part III: Post-experimental questionnaire

After the 12 rounds of the experiment, we conducted several incentivized post-experimental questionnaires that involved testing for: (i) risk aversion test of Holt and Laury (2002); (ii) sincerity test; (iii) social value orientation test; (iv) cognitive reflection test of Frederick (2005), ; (v) social media usage; (vi) demographics.[8]

Procedures

In our six main treatments with information about posteriors we have six independent groups of eight participants per treatment, a total of 288 participants. The experiment was conducted at the LINEEX in Valencia (Spain). Experimental sessions with communication lasted an average of 180 minutes, while the ones without communication had an average of 130 minutes. Furthermore, experimental participants received a fixed participation fee of €5, plus a variable payment based on their performance in 1 of the 12 rounds of the virtual financial market, as well as on their answers to two sets of questions at the end of the experiment. Average earnings were 21 Euros.

3 Results

We first describe the results on market prices and their informativeness in sub-section 3.1, and the ones on liquidity in sub-section 3.2. We then analyze beliefs in sub-section 3.3, and communication in sub-section 3.4.

3.1 Market Prices and Informativeness

Our central measure of market quality is the accuracy of price formation. We assess this through two complementary lenses. First, we analyze the mean transaction price by treatment and fundamental value. This analysis addresses whether prices respond to fundamental information at all. Second, we examine the results of how close transaction prices are to the fundamental value as a measure of price informativeness.

The distribution of transaction prices across treatments is reported in Table 14 of Appendix C, and figures of all transaction prices and realizations of the fundamental value

[8]Details of the post-experimental questionnaires can be found in the Online Appendix.

in each group in each treatment are shown in Section 7.2 of Appendix C.[9]

3.1.1 Market prices and fundamental values

To test whether prices reflect fundamental value, and whether this relationship varies across our experimental treatments, we estimate the following random-effects panel regression:

$$p_{gr} = \alpha + \beta \cdot \text{High}_{gr} + \sum_{k \neq \text{B\_NoCom}} \gamma_k \cdot T_{gk} + \sum_{k \neq \text{B\_NoCom}} \delta_k \cdot (\text{High}_{gr} \times T_{gk}) + u_g + \varepsilon_{gr}, \quad (1)$$

where p_{gr} is the price in group g in round r, High_{gr} is an indicator equal to one when the fundamental value in that round is 500 (and zero when it is 50), and T_{gk} are indicators for the treatment assigned to group g, with B_NoCom (biased information, no communication) as the reference category. The term u_g is a group-level random effect capturing unobserved heterogeneity across player-groups, and ε_{gr} is the idiosyncratic error term. Standard errors are clustered at the player-group level to account for within-group correlation across rounds. The coefficient β captures the high–low price gap in the baseline treatment, while the interaction coefficients δ_k measure how this gap shifts in each alternative treatment. Table 4 reports both the regression coefficients (Panel A) and the marginal effect of moving from low to high fundamental value within each treatment (Panel B).

The most striking feature of the results is the magnitude of mispricing relative to the underlying fundamentals. Given that low-fundamental assets are worth 50 and high-fundamental assets are worth 500. No treatment comes close to this benchmark.

In the no-communication baseline (B_NoCom), the estimated constant of 262.57 indicates that low-fundamental assets (worth 50) trade at roughly five times their fundamental value, while high-fundamental assets (worth 500) trade at approximately 279—only about 56% of their fundamental value. The high–low price difference of 16.16 is small in absolute terms,

[9] Since the experiment imposed no upper bound on bid and ask prices, a small number of transactions were executed at implausibly high prices, likely reflecting data entry errors or irrational behaviour. Specifically, the theoretical maximum fundamental value in the experiment is 500, and we exclude the 45 transactions (0.53% of all trades) with prices exceeding 700 — 40% above this maximum — as these fall well outside any plausible range of rational trading.

Table 4
Effect of High Fundamental Value on Transaction Prices Across Treatments

	Coefficient	**Std. Error**
Panel A: Random-Effects GLS Regression		
High Fundamental (= 1)	16.157	(10.519)
Treatment (ref: B_NoCom)		
B_AllCom	−26.354	(34.862)
B_SegCom	−66.022*	(38.169)
U_AllCom	−52.099	(34.393)
U_NoCom	−52.109	(37.322)
U_SegCom	−109.141***	(40.215)
High Fundamental × Treatment		
High × B_AllCom	18.200	(21.681)
High × B_SegCom	90.787***	(33.396)
High × U_AllCom	83.561***	(27.696)
High × U_NoCom	10.929	(19.886)
High × U_SegCom	55.831**	(25.405)
Constant	262.566***	(32.084)
Panel B: Marginal Effect of High Fundamental by Treatment		
B_NoCom	16.157	(10.519)
B_AllCom	34.358*	(18.958)
B_SegCom	106.944***	(31.696)
U_AllCom	99.718***	(25.620)
U_NoCom	27.086	(16.876)
U_SegCom	71.989***	(23.125)
Observations	432	
Groups (playergroup)	36	
Obs. per group	12	
Wald $\chi^2(11)$	62.70	
Prob > χ^2	0.000	
R^2 (within / between / overall)	0.200 / 0.309 / 0.238	
ρ	0.314	

Notes: Random-effects GLS regression with standard errors clustered at the player-group level (36 clusters) reported in parentheses. The dependent variable is the transaction price. The reference categories are low fundamental value (*high_fundamental* = 0) and treatment B_NoCom. Panel B reports the marginal effect of moving from low to high fundamental value within each treatment.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

represents only 3.6% of the true fundamental gap, and is not statistically distinguishable from zero ($p = 0.125$). In essence, prices in the no-communication treatments are anchored near an intermediate value and largely fail to incorporate the fundamental signal. A similar pattern characterizes U_NoCom, where the marginal effect of high fundamentals is 27.09 (6.0% of the fundamental gap, $p = 0.108$). Removing the channel of communication—regardless of whether the underlying information is biased or unbiased—appears to prevent prices from meaningfully tracking fundamentals.

The treatments that allow communication exhibit substantially larger price responses to fundamentals and the differences are statistically significant in all four of the relevant treatments. The marginal effects in Panel B of Table 4 are 34.36 in B_AllCom ($p = 0.070$), 71.99 in U_SegCom ($p = 0.002$), 99.72 in U_AllCom ($p < 0.001$), and 106.94 in B_SegCom ($p = 0.001$). The interaction terms in Panel A confirm that the high–low price gap is significantly larger than the B_NoCom baseline in B_SegCom ($p = 0.007$), U_AllCom ($p = 0.003$), and U_SegCom ($p = 0.028$).

These effects, though statistically robust, remain economically modest relative to the true fundamental gap. Even in the most responsive treatment (B_SegCom), prices capture only 23.8% of the 450-unit fundamental difference (106.944/450). Decomposing the predicted prices in B_SegCom yields approximately 197 for low-fundamental assets (still nearly four times their fundamental value of 50) and 304 for high-fundamental assets (still 39% below their fundamental value of 500). Communication therefore widens the price gap between high- and low-fundamental assets and makes the difference significant, but does not eliminate the systematic overpricing of low-value assets nor the systematic underpricing of high-value assets.

A second pattern of interest concerns the interaction between the information environment (biased vs. unbiased) and the communication network (no communication, segmented communication, or full communication). The two treatments with the largest price responses are B_SegCom (biased information, segmented communication) and U_AllCom (unbiased information, full communication), suggesting that the most effective communication structure depends on the nature of the underlying information. When signals are biased, segmented communication appears to outperform full communication. When signals are unbiased, by contrast, full communication is most effective, consistent with the intuition that pooling reliable information across the entire market improves aggregation. B_AllCom—the combination of biased information with unrestricted communication—produces the weakest effect among the communication treatments, lending further support to the interpretation that open communication can amplify rather than correct biased signals.

Overall, the results show that communication does produce statistically and economically meaningful improvements in the responsiveness of prices to fundamental value relative to the no-communication baseline. However, even in the most favorable treatment, prices incorporate fundamentals only modestly: low-value assets remain substantially overpriced and high-value assets remain substantially underpriced. Communication is thus best described as necessary but not sufficient for price discovery in this setting. The data further suggest that the most effective communication network depends on the quality of the underlying information, with segmented communication outperforming full communication when signals are biased, and vice versa when signals are unbiased.

3.1.2 Inaccuracy of prices

In this sub-section, the definition of price accuracy is the mean absolute deviation (MAD) between transaction prices and fundamental value defined as:

$$MAD = \frac{1}{t}\sum_{i}(p_i - \theta), \tag{2}$$

where t is the number of transactions in a given market, p_i is the transaction price, and θ is the fundamental value of the asset. Notice, that a higher (lower) MAD implies a lower (higher) price informativeness. Table 5 and Figure 2 report the results.

We then examine the significance of the overall interaction effect on MAD. The joint Wald test for the Info $\times$ NetCom interaction in MAD is highly significant ($\chi^2(5) = 23.82$, $p = 0.0002$), indicating that communication structure and information bias jointly shape the accuracy of price formation in a manner that cannot be reduced to either factor alone.[10]

Figure 2: Mean Absolute Deviation between Price and Fundamental Value

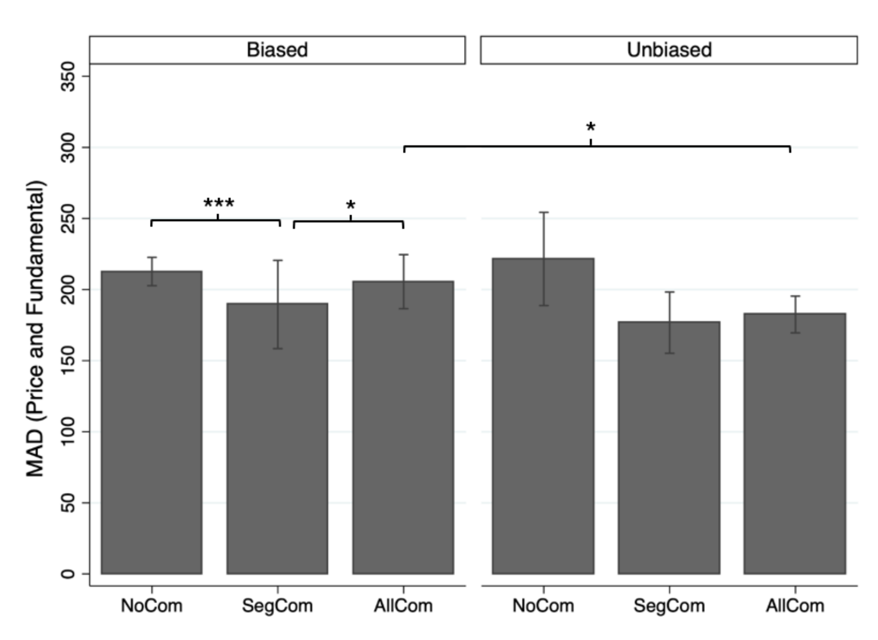

Note: Bars show the average MAD in each treatment. Vertical lines reflect 95% confidence intervals of the average MAD from a random effects panel regression with treatment dummies. The horizontal brackets above the bars indicate pairwise statistical comparisons of the equality of average MAD per treatment using cluster-robust t-tests. Standard errors are clustered at the group level. The asterisks above each bracket indicate the unadjusted significance level for that pairwise comparison: ***, **, * correspond to *p-values* at the 1%, 5%, 10%, respectively. For neat presentation, no stars are reported for treatment differences (where comparison is warranted) that were not significant.

[10] The joint Wald test for the information bias and network structure reported above is highly significant, yet only one pairwise contrast survives Holm–Bonferroni correction. These results are complementary rather than contradictory: the joint test detects a systematic pattern across all six cells simultaneously, while the pairwise comparisons locate its source — SegCom reduces MAD relative to NoCom under biased information but not under unbiased information, suggesting that network structure matters for price accuracy only when the information environment is biased.

Panel A of Table 5 shows that, when information is biased, network structure has a substantial and statistically robust effect on mispricing. Markets without communication exhibit lower price informativeness (higher MAD) than markets with segmented communication, a difference that survives Holm correction at $p = 0.016$ and is the most robust pairwise result in our analysis. Segmented communication, in turn, produces higher price informativeness (lower MAD) than full all-to-all communication ($p = 0.074$ unadjusted; $p = 0.148$ after Holm correction)—a marginal but suggestive reversal. The NoCom–AllCom contrast is small and non-significant ($p = 0.394$).

Table 5

Pairwise Comparisons of MAD across Information Structures and Communication Structure

	Coef.	SE	p (unadj.)	p (Holm)
Panel A: Network structure comparison within Biased Info				
NoCom vs. SegCom	43.22**	15.49	0.005	0.016
NoCom vs. AllCom	9.72	11.42	0.394	0.395
SegCom vs. AllCom	−33.50†	18.74	0.074	0.148
Panel B: Network structure comparison within Unbiased Info				
NoCom vs. SegCom	26.82	20.77	0.197	0.393
NoCom vs. AllCom	28.67	20.86	0.169	0.508
SegCom vs. AllCom	1.84	18.07	0.919	0.919
Panel C: Information bias comparisons within each network structure				
NoCom: Biased vs. Unbiased	11.88	16.72	0.478	0.956
SegCom: Biased vs. Unbiased	−4.52	19.80	0.819	0.819
AllCom: Biased vs. Unbiased	30.82†	16.91	0.068	0.205
Overall interaction: $\chi^2(5) = 23.82$, $p = 0.0002$				

Notes: All contrasts estimated from a single random-effects GLS regression ($N = 432$, 36 groups) with robust standard errors clustered at the group level. Marginal cell means — *Biased Info*: NoCom = 218.11, SegCom = 174.89, AllCom = 208.38. *Unbiased Info*: NoCom = 206.23, SegCom = 179.41, AllCom = 177.56. Holm–Bonferroni correction applied separately within each hypothesis family (Panels A, B, and C). $^{**}p < 0.05$ after Holm correction; $^{\dagger}p < 0.10$ unadjusted only.

Panel B presents a different picture. Under unbiased information, none of the three network structures contrasts is statistically significant: NoCom vs. SegCom ($p = 0.197$), NoCom vs. AllCom ($p = 0.169$), and SegCom vs. AllCom ($p = 0.919$). When the underlying information is unbiased, communication structure does not detectably affect price accuracy.

The asymmetry between Panels A and B is itself a key result: communication structure matters for price accuracy only when information is biased.

Panel C contrasts biased and unbiased information within each network structure. None of the three contrasts reaches conventional significance after Holm correction. The strongest comparison is under AllCom, where biased markets exhibit lower price accuracy (higher MAD) than unbiased markets ($p = 0.068$ unadjusted; $p = 0.205$ after Holm correction). Under NoCom and SegCom, the differences are negligible. Information structure has its largest effect on price accuracy precisely when traders are fully connected.

The two sub-sections 3.1.1 and 3.1.2 therefore tell complementary parts of the same story. Table 4 shows that communication is necessary for prices to respond at all to the fundamental, regardless of bias. Figure 2 and Table 5 show that, conditional on prices responding, segmented communication produces the most accurate prices when information is biased.

Result 1: *i) Communication is necessary for prices to incorporate fundamental information at all—without it, prices barely respond to large differences in underlying value.*[11]

ii) The interaction of the network structure and information type is highly significant for price accuracy, establishing that price formation depends jointly on the information environment and the structure of trader interaction.

iii) Communication structure affects mispricing primarily under biased information: segmented communication produces the most accurate prices, while no-communication and full-communication markets perform comparably and worse. However, Communication does not affect mispricing under unbiased information.

iv) When all participants communicate, biased information leads to more mispricing than unbiased information.

The role of communication networks in price discovery is thus twofold: it is a precondition for prices responding to fundamentals, and—under biased information—a determinant of how

[11] However, low-value assets remain substantially overpriced and high-value assets remain substantially underpriced in all treatments.

closely those prices track them. The liquidity outcomes examined in the following sub-section provide further evidence on the channels through which these effects operate.

3.2 Liquidity

Another dimension of market quality is liquidity. We measure market liquidity using two complementary indicators. First, we use trading volume, which captures the intensity of trading. Second, we employ the relative bid-ask spread as a measure of market liquidity and transaction costs. The bid-ask spread is the difference between the highest price a buyer is willing to pay and the lowest price a seller is willing to accept. A trader who wishes to transact immediately must either buy at the seller's price or sell at the buyer's price, paying this gap as the implicit cost of trading. We therefore use the relative bid-ask spread — the spread expressed as a fraction of the asset's midpoint price - measured at each second of the trading round, and then averaged across all seconds in a round. A narrower spread indicating a more liquid and efficient market.

Tables 15 and 16 in Appendix C show the distribution of actions (trade, withdrawn, and submitted but not executed) across treatments, and the distribution of submitted prices across price bins across treatments, respectively. The relative bid–ask spread is constructed with a three-stage outlier treatment. First, submitted prices exceeding 2,000 ECUs (four times the fundamental value) are excluded prior to any calculation, as these are attributable to data entry errors. Second, the relative spread is winsorised at the 99th percentile within each second of the trading round. Third, the round-average relative spread is winsorised at the 99th percentile at the round level. Seconds with no valid quote on one or both sides are excluded from the round average.

Trading volume, measured as the number of transactions completed per round, averages 21.85 trades (SD = 10.8, range = $[3, 66]$), and the average relative bid-ask spread is 0.59 (SD = 0.3, range = $[0.064, 1.35]$). Table 6 reports pairwise contrasts for trading volume. The overall interaction between information structure and communication structure is marginally significant for trading volume ($\chi^2(5) = 9.58$, $p = 0.088$) and marginally significant for the relative bid–ask spread ($\chi^2(5) = 10.86$, $p = 0.054$), confirming that the effect of network structure on market liquidity depends on information structure, though the two measures tell complementary rather than identical stories.

TABLE 6
PAIRWISE COMPARISONS OF TRADING VOLUME AND RELATIVE BID–ASK SPREAD BY INFORMATION BIAS AND NETWORK COMMUNICATION

	Trading Volume				Relative Bid–Ask Spread			
Comparison	**Coef.**	**SE**	p	p^H	**Coef.**	**SE**	p	p^H
Panel A. NetCom comparisons within Biased Information (Family 1)								
NoCom vs. SegCom	−0.64	3.45	0.853	0.853	−0.2552**	0.0791	0.001	0.004
NoCom vs. AllCom	0.58	2.79	0.834	1.000	−0.1847	0.1127	0.101	0.202
SegCom vs. AllCom	1.22	3.27	0.709	1.000	0.0705	0.1028	0.493	0.493
Panel B. NetCom comparisons within Unbiased Information (Family 2)								
NoCom vs. SegCom	4.71	5.92	0.427	0.427	0.0595	0.1049	0.571	1.000
NoCom vs. AllCom	11.57*	5.58	0.038	0.076	0.0114	0.0979	0.907	1.000
SegCom vs. AllCom	6.86*	3.01	0.023	0.068	−0.0481	0.0870	0.581	1.000
Panel C. Information comparisons within each NetCom level (Family 3)								
NoCom: Biased vs. Unbiased	−6.93	5.75	0.228	0.456	−0.1929†	0.1034	0.062	0.186
SegCom: Biased vs. Unbiased	−1.58	3.73	0.671	0.671	0.1217	0.0810	0.133	0.266
AllCom: Biased vs. Unbiased	4.06†	2.43	0.094	0.283	0.0032	0.1076	0.976	0.976
Overall Info × NetCom interaction (Wald test):								
Trading Volume: $\chi^2(5) = 9.58$, $p = 0.088$ | Bid–Ask Spread: $\chi^2(5) = 10.86$, $p = 0.054$								

Notes. Coefficients reflect contrasts among cell means estimated jointly from random-effects GLS regressions of each outcome on the full Info × NetCom interaction, using $N = 432$ observations from 36 player groups (12 obs. per group), with clustered standard errors at the group level. The relative bid–ask spread is defined as the difference between the best ask and best bid at each second of the trading round, normalised by their midpoint and averaged across all seconds within a round; it is winsorised at the 99th percentile at both the second level and the round-average level, and submitted prices exceeding 2,000 experimental currency units are excluded prior to computation. Standard errors are cluster-robust at the player-group level. Unadjusted p-values (p) are based on Wald z-tests. Holm-adjusted p-values (p^H) are computed within each comparison family (3 tests per family) using the Holm–Bonferroni step-down procedure. The overall Info × NetCom interaction is reported from a joint Wald test of all interaction contrasts. Marginal cell means — *Biased Info*: Trading Volume: NoCom = 19.04, SegCom = 19.68, AllCom = 18.46; Spread: NoCom = 0.430, SegCom = 0.685, AllCom = 0.615. *Unbiased Info*: Trading Volume: NoCom = 25.97, SegCom = 21.26, AllCom = 14.40; Spread: NoCom = 0.622, SegCom = 0.563, AllCom = 0.611. Significance markers (based on unadjusted p): † $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

For trading volume, treatment differences are concentrated entirely within the unbiased information condition (Panel B). Under unbiased information, NoCom and SegCom both generate significantly higher volume than AllCom (NoCom vs. AllCom: coef. = 11.57, p =

0.038; SegCom vs. AllCom: coef. = 6.86, $p = 0.023$), though neither contrast survives Holm correction at the 5respectively). Under biased information, all three network structures generate statistically indistinguishable trading activity (Panel A). A plausible interpretation is that under unbiased information, unrestricted communication allows traders to reach faster agreement on the fundamental value, reducing the incentive to trade in order to discover the fundamental value. Under biased information, by contrast, uncertainty about the true fundamental persists regardless of network structure, sustaining similar levels of trading activity across treatments.

The relative bid–ask spread tells a different story, with treatment differences concentrated within the biased information condition (Panel A). Under biased information, NoCom exhibits a significantly narrower spread than SegCom (coef. = −0.255, $p_{\text{Holm}} = 0.004$), indicating that segmented communication is associated with wider spreads and therefore higher transaction costs relative to the no-communication baseline. The NoCom–AllCom contrast points in the same direction but is not significant (coef. = −0.185, $p = 0.101$), while SegCom and AllCom do not differ significantly (coef. = 0.071, $p = 0.493$). Under unbiased information, by contrast, spread differences across all three network structures are small and entirely non-significant (Panel B), suggesting that communication structure has no meaningful effect on transaction costs when information is accurate. This asymmetry mirrors the pattern observed for price accuracy and belief formation: network structure shapes market outcomes primarily when the information environment is biased.

Panel C confirms that holding network structure fixed, information structure has limited direct effects on either liquidity measure. The only marginally significant contrast is for the spread under NoCom (coef. = −0.193, $p = 0.062$), where biased traders face a somewhat wider spread than unbiased traders, though this does not survive Holm correction ($p^H = 0.186$). For trading volume, the AllCom comparison is marginally significant in the unadjusted test (coef. = 4.06, $p = 0.094$), with biased AllCom markets generating slightly more volume than unbiased AllCom markets, but again this does not survive correction ($p^H = 0.283$).

The results can be summarized as follows

Result 2:

i) [Trading Volume] Communication structure affects trading volume only under unbiased information. When information is unbiased, communicating with the whole group (AllCom) significantly reduces trading volume relative to both no communication (NoCom) and segmented communication (SegCom). Under biased information, trading volume is statistically indistinguishable across all three network structures.

ii) [Relative spread] Communication structure affects liquidity asymmetrically across information environments. Under unbiased information, the relative bid–ask spread does not differ significantly across network structures. While for the biased information structure, the relative bid-ask spread is higher for segmented communication (SegCom) comapred to no communication (NoCom).

3.3 Beliefs

Beliefs were elicited before trading and after communication, providing insight into participants' accuracy and reasoning processes in identifying the color of the ball (i.e., high or low fundamental) prior to the trading phase.[12]

We measure the pre-trade belief accuracy as the mean absolute deviation between each trader's stated belief and the color of the ball (which represents the asset's fundamental value). Recall that beliefs were elicited in all treatments after the information (and communication) stages. Figure 3 displays the marginal cell means across all six treatment conditions, and Table 7 reports the pairwise contrasts with Holm–Bonferroni corrected p-values.

The overall interaction between information structure and communication structure is statistically significant ($\chi^2(5) = 30.83$, $p < 0.0001$), indicating that the effect of network structure on belief accuracy depends on whether traders receive biased or unbiased information and on

[12]We also elicited beliefs after trading. The essence of the results reported in this section hold. See Table 18 in the Appendix.

the communication network. We therefore examine each information condition separately.

Figure 3: Group Level Mean Absolute Deviation between Pre-trade Belief and the Color of the Ball

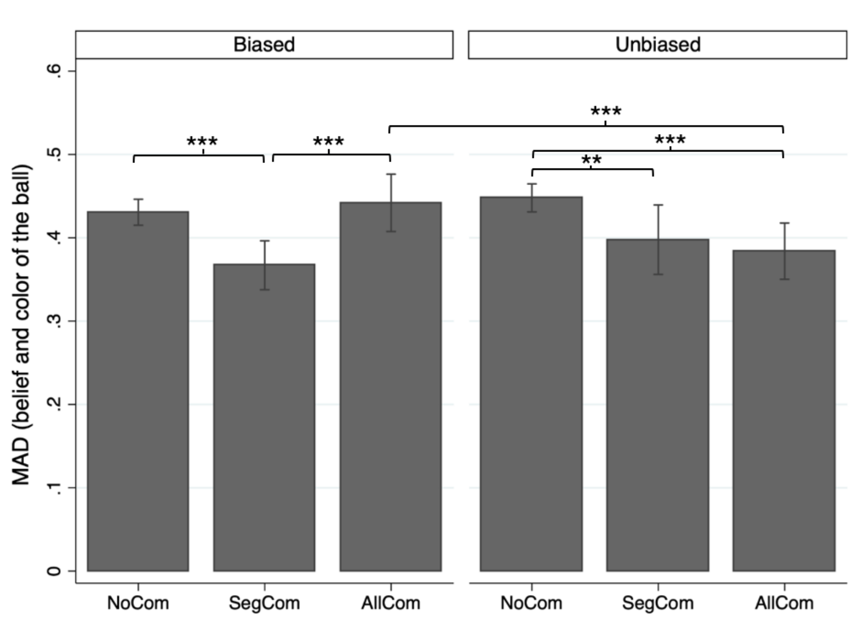

Note: Bars show the average mean absolute deviation between the belief and color of the ball in each treatment. Vertical lines reflect 95% confidence intervals of the average average mean absolute deviation between the belief and color of the ball from a random effects panel regression with treatment dummies. The horizontal brackets above the bars indicate pairwise statistical comparisons of the equality of average mean absolute deviation between the belief and color of the ball per treatment using cluster-robust t-tests. Standard errors are clustered at the group level. The asterisks above each bracket indicate the significance level for that pairwise comparison: ***, **, * correspond to *p-values* at the 1%, 5%, 10%, respectively. For neat presentation, no stars are reported for treatment differences (where comparison is warranted) that were not significant.

Under biased information, network structure has a pronounced effect on belief accuracy. Traders in the segmented communication condition (SegCom) hold substantially more accurate beliefs than those in either the no-communication baseline or the all communication condition. The contrast between NoCom and SegCom is large and survives Holm correction (coef. = 0.064, $p_{\text{Holm}} = 0.001$), as does the contrast between SegCom and AllCom (coef. = -0.075, $p_{\text{Holm}} = 0.002$). Notably, opening communication to all traders (AllCom) does not improve on the no-communication baseline — the two conditions are statistically indistinguishable (coef. = -0.011, $p = 0.557$) and nearly identical in magnitude. This suggests that under biased information, only a restricted communication structure improves belief accuracy, while unrestricted communication appears to impair it.

The pattern under unbiased information is markedly different and, in some respects, the mirror image of the biased case. Here it is AllCom that produces the most accurate beliefs, followed by SegCom, with NoCom yielding the least accurate beliefs. Both contrasts involving NoCom are statistically significant after Holm correction: NoCom vs. SegCom (coef. = 0.050, $p_{\text{Holm}} = 0.057$, marginal) and NoCom vs. AllCom (coef. = 0.064, $p_{\text{Holm}} = 0.003$). The difference between SegCom and AllCom is small and non-significant (coef. = 0.014, $p = 0.615$). Under unbiased information, therefore, more communication is unambiguously beneficial — broader information sharing helps traders converge toward the true fundamental value.

Holding communication structure fixed, the direct comparison between biased and unbiased information reveals a significant difference only in the AllCom condition (coef. = 0.058, $p_{\text{Holm}} = 0.055$), where biased traders hold less accurate beliefs than their unbiased counterparts. In NoCom and SegCom, the two information conditions produce similar levels of belief accuracy, suggesting that when communication is absent or restricted, the quality of the initial signal matters less for belief formation.[13]

Taken together, these results point to an important interaction between information structure and communication structure in determining belief accuracy. Restricted communication (SegCom) is most beneficial when information is biased. Conversely, communicating with the whole group (AllCom) is most beneficial when information is unbiased, as it allows

[13] In Table 18 in the Appendix, we report the results for belief dispersion among the traders in a group. We find that the communication and information structures do not shape the degree of disagreement among traders within a group.

Table 7
Pairwise Comparisons of Mean Absolute Deviation between Beliefs and Color of the Ball across Information and Communication Structures

	Coef.	SE	p (unadj.)	p (Holm)
Panel A: Network structure comparison within Biased Info				
NoCom vs. SegCom	0.0636**	0.0170	0.000	0.001
NoCom vs. AllCom	−0.0113	0.0193	0.557	0.557
SegCom vs. AllCom	−0.0749**	0.0231	0.001	0.002
Panel B: Network structure comparison within Unbiased Info				
NoCom vs. SegCom	0.0502	0.0229	0.029	0.057
NoCom vs. AllCom	0.0639**	0.0192	0.001	0.003
SegCom vs. AllCom	0.0138	0.0274	0.615	0.615
Panel C: Information bias comparisons within each network structure				
NoCom: Biased vs. Unbiased	−0.0172	0.0117	0.140	0.281
SegCom: Biased vs. Unbiased	−0.0307	0.0260	0.238	0.238
AllCom: Biased vs. Unbiased	0.0580†	0.0246	0.018	0.055
Overall interaction: $\chi^2(5) = 30.83$, $p < 0.0001$				

Notes: All contrasts estimated from a single random-effects GLS regression ($N = 432$, 36 groups) with robust standard errors clustered at the group level. Holm–Bonferroni correction applied separately within each hypothesis family (Panels A, B, and C). $^{**}p < 0.05$ after Holm correction; $^{\dagger}p < 0.10$ unadjusted only.

accurate signals to diffuse widely. Unrestricted communication under biased information appears to be the worst of both worlds, producing belief accuracy no better than complete communication silence.

Result 3:

i) Pre-trade belief accuracy depends critically on both information and network structures.

ii) With biased signals, segmented communication yields the most accurate beliefs, outperforming both no communication and full communication.

iii) With unbiased signals any form of communication (segmented or full) improves belief accuracy relative to no communication.

iv) Full communication performs best when information is unbiased.

Note that these belief accuracy patterns mirror the results of price informativeness.

3.4 Communication

In our experiment, communication was free and unrestricted and it provides a rich dataset that can offers a window into participants' thinking processes. In this sub-section, we focus on groups in the communication treatments (B_SegCom, B_AllCom, U_SegCom and U_AllCom). We used two human coders to analyze the chat data at the group level for each round. See Appendix B.1 for the coder's protocol, Appendix B.2 for the analysis of the consistency among coders' answers in the variables which reached agreement (Cohen's kappa above 0.55).

In this sub-section, all treatment comparisons are estimated from a single random-effects GLS regression of the outcome on the full set of treatment cell indicators at the group-round level ($N = 432$, 36 groups), with robust standard errors clustered at the group level. Pairwise contrasts are applied to the estimated cell means. The overall interaction between information structure and communication structure is assessed with a single joint Wald test, requiring no multiplicity correction. The results are presented in Table 8.

Analysis of the results show that across all treatments, an average of 64.5% of participants say something in the chat, and 54% share relevant information (either their signal or their belief about the color of the ball). Differences across treatments in these dimensions are not statistically significant (number of participants chatting: $\chi^2(2) = 0.21$, $p = 0.899$; share of relevant information shared: $\chi^2(2) = 2.32$, $p = 0.314$).

However, the number of messages per participant per round differs across treatments ($\chi^2(3) = 7.50$, $p = 0.058$), driven primarily by participants in the U_SegCom condition sending significantly more messages than those in U_AllCom (coef. = 0.605, $p = 0.006$), while no significant difference is found between B_SegCom and B_AllCom (coef. = 0.048, $p = 0.914$). Additionally, we find the following regarding false information sharing: the interaction between information structure and communication structure is significant ($\chi^2(2) = 6.07$, $p = 0.048$), with participants in U_SegCom sharing significantly less false information than those in U_AllCom (coef. = -1.410, $p < 0.001$), whereas no such difference emerges under biased information (coef. $= -0.438$, $p = 0.250$). The direct comparison between biased and unbiased traders within AllCom is also non-significant for both outcomes ($p = 0.220$ and $p = 0.288$, respectively).

We next study how communication translates into shared beliefs by examining whether groups reach a majority consensus on a common view of the asset's fundamental value. To do so, we classify each group-round into one of three outcomes: no consensus or consensus (if the majority of a communication group agree on the color of the ball), which is split between false consensus or true consensus. We then compare the incidence of these outcomes across treatments to assess how information structure and network structure shape collective agreement. The graphical results are presented in Figure 4, and the statistical tests in Panel B of Table 8.

Figure 4: Consensus about the color of the ball: prevalence, true or false.

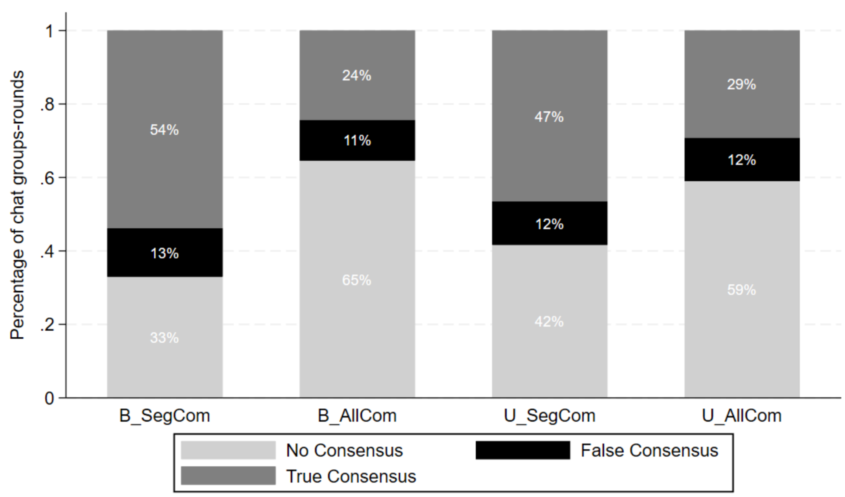

Note: This graph shows the % of chat-group-rounds that we classified according to the average of two coders in one of the three categories: (i) did not reach a consensus (No Consensus); (ii) reached a consensus but on the wrong color of the ball (False Consensus); and (iii) reached a consensus on the right color of the ball (True Consensus) in each treatment.

For overall consensus ($\chi^2(3) = 12.05$, $p = 0.007$), SegCom generates significantly higher consensus than AllCom under both biased (coef. $= 0.316$, $p = 0.006$) and unbiased information

Table 8
Pairwise Comparisons by Family: Communication Network and Information Type Effects on Chat and Consensus Outcomes

Family	Comparison	Panel A: Chat Outcomes			
		(1) No. Part.	(2) Share Info	(3) Mess./Part.	(4) False Info
Interaction test: joint Wald $\chi^2(2)$, p-value					
	Biased & Unbiased × NetCom	0.899	0.314	0.420	0.048**
Family 1 — Network communication within Biased information					
	SegCom vs. AllCom	0.007	0.833	0.048	−0.438
		(0.976)	(0.235)	(0.914)	(0.250)
Family 2 — Network communication within Unbiased information					
	SegCom vs. AllCom	0.063	−0.597	0.605***	−1.410***
		(0.747)	(0.357)	(0.006)	(< 0.001)
Family 3 — Information type, holding network communication fixed					
	AllCom: Biased vs. Unbiased	0.069	−0.632	0.366	−0.354
		(0.647)	(0.266)	(0.220)	(0.288)
Observations		288	288	288	288
Groups		24	24	24	24

Family	Comparison	Panel B: Consensus Outcomes		
		(5) Consensus	(6) False Cons.	(7) True Cons.
Interaction test: joint Wald $\chi^2(2)$, p-value				
	Biased & Unbiased × NetCom	0.563	0.946	0.598
Family 1 — Network communication within Biased information				
	SegCom vs. AllCom	0.316***	0.021	0.295***
		(0.006)	(0.745)	(0.005)
Family 2 — Network communication within Unbiased information				
	SegCom vs. AllCom	0.174†	0.000	0.174**
		(0.057)	(1.000)	(0.032)
Family 3 — Information type, holding network communication fixed				
	AllCom: Biased vs. Unbiased	−0.056	−0.007	−0.049
		(0.631)	(0.918)	(0.642)
Observations		288	288	288
Groups		24	24	24

Notes: Coefficients from `lincom` post-estimation; p-values in parentheses. All regressions are random-effects GLS with standard errors clustered at the group level. The interaction test row reports the p-value of the joint Wald $\chi^2(2)$ test for the Info × NetCom interaction. Pairwise comparisons within each family are subject to Holm–Bonferroni correction. Panel A dependent variables: (1) number of chat participants; (2) share of informative messages; (3) messages per participant per round; (4) share of false information. Panel B dependent variables: (5) average consensus on ball colour; (6) false consensus; (7) true consensus.
$^{\dagger}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

(coef. = 0.174, p = 0.057, marginal), while the interaction between information structure and communication structure is not significant ($\chi^2(2) = 1.15$, $p = 0.563$). A similar pattern emerges for true consensus ($\chi^2(3) = 14.69$, $p = 0.002$): SegCom produces significantly more true consensus than AllCom under both biased (coef. = 0.295, p = 0.005) and unbiased information (coef. = 0.174, p = 0.032), with no significant difference between information conditions within AllCom (coef. = −0.049, p = 0.642) and no significant interaction ($\chi^2(2) = 1.03$, $p = 0.598$).

By contrast, false consensus shows no meaningful variation across treatments. The overall interaction is far from significant ($\chi^2(2) = 0.11$, $p = 0.946$), and none of the pairwise contrasts approach conventional significance levels — neither SegCom vs. AllCom within biased (coef. = 0.021, $p = 0.745$) or unbiased information (coef. = 0.000, $p = 1.000$), nor the comparison between information conditions within AllCom (coef. = −0.007, $p = 0.918$).

Result 4: *SegCom consistently achieves higher overall consensus than AllCom regardless of information structure, and this advantage is driven entirely by true consensus: SegCom generates significantly more true agreement than AllCom under both biased and unbiased information. This pattern is consistent with the finding that participants in U_SegCom send more messages and share significantly less false information than those in U_AllCom. By contrast, false consensus rates are statistically indistinguishable across all treatments.*

These results suggest that the communication advantage of segmented networks operates through more active and accurate information sharing rather than through any differential tendency to generate false agreement.

4 Mechanisms

In this section, we examine the mechanisms behind our treatment effects. Subsection 4.1 relates price informativeness to belief accuracy; subsection 4.2 turns to the communication treatments, analyzing the chat data to study how belief accuracy depends on chat participation and agreement.

4.1 Price Informativeness and Beliefs

We begin by examining the impact of the experimental conditions on market mispricing, measured by the mean absolute deviation (MAD) of transaction prices from the fundamental value. Table 9 reports random-effects GLS regressions of MAD on indicators for each treatment cell, with Biased × NoCom as the omitted reference category. All specifications cluster standard errors at the group level and include a linear Round trend to capture within-session learning.

Column (1) presents the baseline specification with treatment indicators and the round trend only. Column (2) adds two controls: an indicator for high-fundamental rounds and an indicator for signals with higher posterior informativeness. Column (3) introduces two belief measures—the mean absolute belief error and the within-period standard deviation of beliefs—to assess the extent to which treatment effects operate through the belief-formation channel. Column (4) instead introduces two market microstructure variables, trading volume and the average relative bid–ask spread, to assess a potential mediation channel. Column (5) combines all controls. Because the belief and microstructure measures are themselves outcomes of the experimental conditions, the coefficients in columns (3)–(5) should not be interpreted as causal treatment effects but rather as direct effects conditional on these intermediate outcomes.

The indicator for high-fundamental rounds enters strongly and positively across columns (2)–(5) (coefficient ≈ 69–73, $p < 0.001$), reflecting the mechanical scaling of MAD with the level of the fundamental. The signal more informative indicator enters at $\beta = -62.0$ in column (2) ($p < 0.001$): more informative signals exhibited substantially higher price informativeness (lower MAD). This coefficient attenuates to -18.5 in column (3) and to -19.0 in column (5) and loses statistical significance, suggesting that the effect of signal informativeness on mispricing operates almost entirely through improved belief accuracy rather than through any independent channel.

Columns (3)–(5) speak to the mechanisms through which the treatments affect mispricing.

Table 9
Random-Effects Panel Regressions of Mis-pricing (MAD Price and Fundamental)

	(1)	(2)	(3)	(4)	(5)
Treatments (omitted: Biased × NoCom)					
Biased × SegCom	−43.220***	−39.260**	−30.692**	−50.746***	−41.509**
	(15.512)	(15.347)	(14.209)	(18.373)	(17.567)
Biased × AllCom	−9.723	−14.187	−14.465	−23.068**	−22.113*
	(11.432)	(10.980)	(11.378)	(10.813)	(12.019)
Unbiased × NoCom	−11.879	−5.298	−5.286	−10.455	−9.268
	(16.744)	(14.001)	(13.711)	(13.452)	(12.503)
Unbiased × SegCom	−38.700***	−35.736***	−25.473**	−40.654***	−30.085**
	(13.101)	(12.637)	(12.132)	(11.948)	(12.249)
Unbiased × AllCom	−40.544***	−32.926***	−26.384***	−43.847***	−36.742***
	(13.238)	(11.880)	(9.996)	(11.200)	(9.883)
Round and information controls					
Round	−2.397	−1.777	−1.809	−1.401	−1.502
	(1.458)	(1.350)	(1.316)	(1.411)	(1.403)
High fundamental		69.031***	72.466***	69.692***	72.870***
		(17.955)	(17.911)	(17.764)	(17.736)
Signal more info		−62.027***	−18.491	−61.262***	−19.047
		(13.985)	(15.277)	(13.120)	(15.288)
Beliefs					
Mean absolute belief error			161.719***		158.427***
			(22.202)		(22.133)
SD of beliefs			−30.468		−45.785
			(50.548)		(51.593)
Market microstructure					
Trading volume				−0.539	−0.558
				(0.516)	(0.493)
Average relative spread				46.325***	41.605**
				(15.948)	(16.225)
Constant	233.688***	227.582***	136.617***	214.708***	131.934***
	(10.884)	(16.713)	(17.089)	(20.743)	(23.815)
Observations	432	432	432	432	432
Groups	36	36	36	36	36
R^2 (overall)	0.043	0.221	0.275	0.245	0.295
R^2 (within)	0.009	0.193	0.247	0.207	0.259
R^2 (between)	0.244	0.400	0.453	0.472	0.511
Wald χ^2	25.99	67.19	166.68	74.02	257.10

Notes: Random-effects GLS regressions of market mispricing, measured as the mean absolute deviation (MAD) of transaction prices from the fundamental value, on treatment indicators and controls. Standard errors (in parentheses) are clustered at the group level (36 clusters). The omitted treatment is *Biased* × *NoCom*. The full set of six treatments is included; only five non-baseline cells appear because the sixth is the omitted category. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The mean absolute belief error enters with a large, precisely estimated coefficient in both column (3) ($\beta = 161.7$, $p < 0.001$) and column (5) ($\beta = 158.4$, $p < 0.001$), confirming that belief accuracy is a primary proximate driver of mispricing. The introduction of the belief-error control in column (3) attenuates the significant treatment coefficients by roughly 20–30% relative to column (2). Taken at face value, this attenuation suggests that roughly 20–30% of the treatment effect on mispricing operates through belief inaccuracy, with the remaining share operating through other channels. We treat this decomposition as suggestive rather than definitive, since it relies on belief accuracy being unconfounded with other unobserved drivers of trading.

Column (4) examines whether the remaining effect operates through market microstructure. The average relative bid–ask spread enters positively and significantly ($\beta = 46.3$, $p = 0.004$), indicating that wider spreads are associated with greater mispricing, while trading volume itself does not significantly predict MAD. The treatment coefficients in column (4) are if anything slightly larger in magnitude than those in column (2)—Biased × SegCom moves from -39.3 to -50.7, and Biased × AllCom becomes significant at -23.1 ($p = 0.033$). Microstructure variables therefore do not mediate the treatment effects; if anything, controlling for them sharpens the estimated contrasts.

Column (5) jointly controls for both belief and microstructure variables and clarifies the relative contribution of each channel. The mean absolute belief error remains the dominant predictor ($\beta = 158.4$, $p < 0.001$), with a coefficient nearly identical to that in column (3). The average relative spread remains significant ($\beta = 41.6$, $p = 0.010$) but slightly attenuated relative to column (4), consistent with a modest overlap between the two channels. The treatment coefficients in column (5) fall between those of columns (3) and (4). All three remain significant at conventional levels, and Biased × AllCom is marginally so (-22.1, $p = 0.066$). Across all five specifications, the rank ordering of treatment effects is preserved and the unbiased treatments together with segmented communication under bias consistently emerge as those most strongly associated with reduced mispricing. Taken together, columns (3)–(5) point to the belief-formation channel as the primary identifiable mechanism behind the treatment effects, with market microstructure playing a separate, non-mediating role.

4.2 Beliefs and Communication

To further analyze the role and determinants of beliefs, Table 10 reports random-effects GLS regressions of the mean absolute belief error on treatment indicators and a sequence of controls, restricting attention to the four communication treatments (Biased and Unbiased crossed with All_Com and Seg_Com). The omitted category is Biased $\times$ AllCom, so each treatment coefficient measures the difference in average absolute belief error relative to subjects who received Biased information and could communicate with the entire group. Specification (1) reports the raw treatment differences with a control for round; specification (2) adds two information controls (whether the realized fundamental is high, and whether the subject's private signal is a more informative one); specification (3) layers on participation measures from the chat data; specification (4) instead adds two measures of the consensus formed during chat — the average share of group members converging on a false and a true state, respectively; and specification (5) includes all controls jointly. Standard errors are clustered at the group level throughout, giving 24 clusters in every specification.

Three patterns emerge from Table 10. Raw treatment differences are statistically meaningful but largely absorbed by information and consensus controls. Without controls (column 1), all three non-baseline cells show lower mean absolute belief errors than the *Biased* $\times$ *AllCom* baseline. The difference is largest and most precisely estimated for *Biased* $\times$ *SegCom* (-0.075, $p < 0.01$), suggesting that segmenting communication is particularly helpful when the underlying information environment is biased. *Unbiased* $\times$ *AllCom* also has a significantly lower mean absolute belief error (-0.058, $p < 0.05$). Once round, fundamental, and signal-structure controls are added (columns 2 and 3), the treatment effects attenuate by roughly one-third, indicating that part of the raw difference is mechanically driven by the information structure of the rounds rather than by the communication protocol itself. After accounting for chat consensus (columns 4 and 5), only *Unbiased* $\times$ *AllCom* remains statistically distinguishable from the baseline (-0.042, $p < 0.01$ in the full specification), while the *Biased* $\times$ *SegCom* effect drops to -0.021 and loses significance. The implication is that the effect of segmented communication under bias operates through the consensus that emerges in chat, rather than being an independent channel.

For the information controls, subjects who receive the more informative private signal have a much lower mean absolute belief error—the coefficient on *Signal more info* is highly significant in every specification ($p < 0.01$). The *High fundamental* indicator carries a small negative coefficient that becomes significant once consensus controls are added, consistent with high-fundamental rounds being marginally easier to read.

The consensus channel does most of the explanatory work. Adding the two consensus measures (column 4) raises overall R^2 from 0.45 to 0.58, with comparable gains in the within and between dimensions. The two consensus variables enter with opposing signs of similar magnitude: agreement on the *False consensus* raises the mean absolute belief error by roughly 0.13–0.15 per unit, while agreement on the *True consensus* lowers it by approximately the same amount, both significant at $p < 0.01$. By contrast, the participation measures (column 3) explain little once consensus is held fixed: of the four chat variables, only *Chat: share of false info* is significant on its own (column 3, 0.033, $p < 0.01$), and even that coefficient drops by three-quarters and loses significance once consensus is included (column 5). Taken together, this points to belief accuracy being driven less by how much chatting occurs than by the content of what subjects end up believing each other agreed on.

The variance decomposition is consistent with this reading. The intraclass correlation ρ falls from 0.077 in column 2 to effectively zero in column 5, meaning that once observable round-level features and chat consensus are accounted for, almost all of the remaining variation in belief error is idiosyncratic rather than group-specific. In other words, the consensus measures account for most of the across-group heterogeneity that the random effects were capturing.

TABLE 10
RANDOM-EFFECTS PANEL REGRESSIONS OF MEAN ABSOLUTE BELIEF ERROR

	(1)	(2)	(3)	(4)	(5)
Treatments (omitted: Biased × AllCom)					
Biased × SegCom	-0.0749^{***}	-0.0513^{**}	-0.0327^{*}	-0.0246	-0.0211
	(0.0233)	(0.0218)	(0.0177)	(0.0162)	(0.0153)
Unbiased × AllCom	-0.0580^{**}	-0.0362	-0.0459^{***}	-0.0366^{**}	-0.0418^{***}
	(0.0248)	(0.0226)	(0.0170)	(0.0168)	(0.0160)
Unbiased × SegCom	-0.0442	-0.0472	-0.0111	-0.0231	-0.0110
	(0.0278)	(0.0290)	(0.0276)	(0.0198)	(0.0203)
Round and information controls					
Round	-0.0005	0.0003	0.0005	-0.0001	0.0011
	(0.0032)	(0.0024)	(0.0027)	(0.0021)	(0.0024)
High fundamental		-0.0216^{*}	-0.0214^{*}	-0.0324^{***}	-0.0302^{**}
		(0.0127)	(0.0125)	(0.0122)	(0.0123)
Signal more info		-0.3612^{***}	-0.3540^{***}	-0.3001^{***}	-0.2986^{***}
		(0.0270)	(0.0274)	(0.0286)	(0.0307)
Participation (chat)					
Chat: # participants			0.0141		0.0212
			(0.0167)		(0.0136)
Chat: share of info			-0.0056		0.0035
			(0.0054)		(0.0043)
Chat: messages / round			-0.0061		-0.0058
			(0.0139)		(0.0104)
Chat: share of false info			0.0325^{***}		0.0082
			(0.0076)		(0.0055)
Consensus					
False consensus (avg.)				0.1493^{***}	0.1331^{***}
				(0.0274)	(0.0266)
True consensus (avg.)				-0.1182^{***}	-0.1256^{***}
				(0.0289)	(0.0312)
Constant	0.4452^{***}	0.6459^{***}	0.5345^{***}	0.6339^{***}	0.4467^{***}
	(0.0234)	(0.0270)	(0.1134)	(0.0208)	(0.0913)
Observations	288	288	288	288	288
Groups	24	24	24	24	24
R^2 (overall)	0.034	0.401	0.450	0.580	0.592
R^2 (within)	0.000	0.409	0.439	0.568	0.571
R^2 (between)	0.295	0.345	0.535	0.675	0.754
Wald χ^2	11.35	230.24	310.18	435.36	645.68

Notes: Random-effects GLS regressions of mean absolute belief error on treatment indicators and controls. Standard errors (in parentheses) are clustered at the group level (24 clusters). The omitted treatment is *Biased × AllCom*. Sample restricted to the four treatments *Biased×AllCom*, *Biased×SegCom*, *Unbiased×AllCom*, and *Unbiased×SegCom*. $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

5 Robustness: the role of posteriors

A natural question is whether our results are driven by the fact that, in each round, subjects were provided with the posterior probability of the ball's color given the prior, the advisor, and the signal received. To test the robustness of our findings to this design choice, we conducted three additional sessions of each of the segmented communication treatments (B_SegCom and U_SegCom) in which the posterior probability was withheld, both from the information screen and from the instructions (see the Online Appendix). These sessions had 24 participants each, with a total of 48 participants. These sessions were otherwise identical to the original treatments except from the number of rounds (8 in the robustness sessions vs. 12 in the main treatments), and number of trial rounds (1 in robustness treatments and 2 in the main treatments). These changes were done for practical purposes.

Table 11 reports the effects of posterior information and information structure on price accuracy, trading volume, and belief accuracy. None of the coefficients on posterior information, information structure, or their interaction reaches conventional levels of statistical significance in any of the three regressions. The overall Wald tests are non-significant for price accuracy ($\chi^2(3) = 3.10$, $p = 0.377$) and belief accuracy ($\chi^2(3) = 3.29$, $p = 0.349$), and marginally significant for trading volume ($\chi^2(3) = 10.20$, $p = 0.017$), though no individual coefficient in that regression is significant. We therefore find no evidence that the availability of posterior information systematically affects market outcomes or belief formation.

6 Discussion

Our experiment yields four interrelated findings on how the structure of communication networks and the information structure jointly shape market outcomes. We discuss their implications and the mechanisms they reveal.

First, across all six treatments, transaction prices deviate substantially from fundamental values: low-value assets are systematically overpriced and high-value assets systematically

Table 11
Random-Effects Panel Regressions: Effects of Posterior Information and Information Type

	(1) MAD Price	(2) Trading Volume	(3) MAD Belief
Posterior	−40.779	4.410	−0.043
	(25.057)	(2.950)	(0.034)
Unbiased Info	−36.426	−3.653	−0.003
	(26.770)	(2.500)	(0.039)
Posterior × Unbiased Info	40.947	5.236	0.033
	(33.387)	(4.516)	(0.047)
Constant	215.665***	15.271***	0.410***
	(19.847)	(1.083)	(0.031)
Observations	238	238	192
Groups	24	24	18
R^2 (overall)	0.019	0.170	0.014
R^2 (within)	0.000	0.000	0.000
R^2 (between)	0.138	0.332	0.136
Wald χ^2	3.10	10.20	3.29
p-value	0.377	0.017	0.349

Notes. Random-effects GLS regressions with robust standard errors clustered at the group level (in parentheses). The dependent variables are: (1) mean absolute deviation (MAD) between transaction prices and the fundamental value; (2) trading volume (number of transactions per round); (3) mean absolute deviation between pre-trade beliefs and the fundamental value. *Posterior* is an indicator equal to one for rounds in which traders receive posterior information. *Unbiased Info* is an indicator equal to one for the unbiased information treatment. *Posterior* × *Unbiased Info* is the interaction between the two. The omitted category is biased information without posterior updating. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

underpriced, even in our most responsive treatment. Communication markedly improves price responsiveness relative to no-communication baselines, but it does not eliminate mispricing. This qualified result echoes a recurring finding in experimental asset markets—that institutional features can move prices toward fundamentals without delivering full efficiency (Plott and Sunder, 1988; Corgnet *et al.*, 2022).

Second, a central contribution of our design is the joint manipulation of information (biased vs. unbiased) and network structure (NoCom, SegCom, AllCom). The data reveal a robust interaction: when information is biased, segmented communication produces the most accurate prices and beliefs and the strongest price response to fundamentals; when information is unbiased, full communication is most effective. Unrestricted communication under biased information emerges as the worst of both worlds, generating belief accuracy no better than complete silence.

Our interpretation rests on a trade-off between two forces. Full communication pools the largest number of signals, which in principle offers the best chance of recovering the fundamental — and indeed, when signals are unbiased, AllCom aggregates information most efficiently and discovers the fundamental value most accurately. Segregation, by contrast, sacrifices breadth of aggregation but makes it easier for a group to reach a strong shared view. When signals are biased, this second force dominates: in AllCom, traders draw on two oppositely biased sources, and the large, mixed conversation is difficult to interpret, so prices fail to reflect fundamentals; in SegCom, each island is smaller and informed by a single source, which makes the fundamental easier to unravel and allows true consensus to form. This implies that expanding communication — through social media or trader chat rooms — does not necessarily move prices closer to fundamentals: when many traders share signals carrying different biases, the resulting conversation may be hard to interpret, consistent with the social transmission bias mechanism of Hirshleifer (2020).

Third, the mechanism analysis in Section 4 traces the treatment effects through belief accuracy and market microstructure. Adding the absolute belief error to the mispricing regression attenuates the treatment coefficients by 20–30%, while the mean absolute belief error itself enters with a large and precisely estimated coefficient. Microstructure variables do not mediate the treatment effects. The chat analysis reinforces this conclusion: once we condition on whether groups reach true or false consensus, the treatment effect of segmented communication under bias loses significance, indicating that segmented networks improve price accuracy primarily by enabling traders to converge on accurate shared beliefs. The advantage of SegCom over AllCom in producing *true* consensus—without any difference in *false* consensus—further suggests that restricted networks operate by promoting accurate information aggregation rather than by magnifying erroneous agreements. This finding complements the experimental evidence of Corgnet *et al.* (2024), who show that communication improves market's informational efficiency, and our results extend it by demonstrating that the structure of the communication network, not merely its presence, is the decisive factor when information is biased.

Fourth, network structure affects the two liquidity measures in opposite information environments. Under unbiased information, full communication reduces trading volume relative to both segmented and no-communication treatments. Under biased information, segmented communication widens bid–ask spreads relative to no communication, indicating higher transaction costs. The two measures therefore capture distinct phenomena—intensity of trading and cost of trading, hence liquidity is not a unitary construct in our setting.

Summarizing, the results offer a coherent message for the design of information environments. Communication structure is not neutral with respect to market quality: its effects depend on the distortions of the signals it transmits. In fully connected networks, unbiased signals aggregate efficiently, while biased signals impede rather than promote information aggregation. Segmented networks may serve as a useful institutional response to biased information, while complete networks are best suited to environments in which underlying signals are unbiased. The findings also suggest that interventions aimed at improving market efficiency should target the belief-formation process—in particular the formation of accurate shared understanding—rather than acting on trading activity or transaction costs directly.

7 Conclusions

This paper examines how communication network structure and information distortion jointly shape market outcomes in a laboratory asset market. Using a 2×3 between-subject design, we vary whether traders receive biased or unbiased private signals and whether they communicate through a fully connected network, two segregated islands, or not at all. Our central finding is that communication is a necessary condition for price discovery — without it, transaction prices fail to differentiate between assets of substantially different fundamental value — but that the effects of the network structure depend critically on the quality of the underlying information. When signals are unbiased, full communication produces the most accurate prices and beliefs. When signals are biased, segmented communication outperforms both full communication and silence across all dimensions of market quality we examine: price accuracy, belief accuracy, and true consensus formation. Unrestricted communication under biased information produces outcomes no better than complete silence.

The mechanism analysis traces these effects through the belief-formation channel. Adding belief accuracy to our mispricing regressions attenuates treatment coefficients by 20–30%, and the advantage of segmented communication under biased information disappears once we condition on the consensus formed during chat, indicating that restricted networks improve market quality primarily by enabling traders to converge on accurate shared beliefs rather than through any independent trading channel. Segmented networks achieve more true consensus — without any difference in false consensus — suggesting they operate by promoting accurate information aggregation. Liquidity results show that under biased information, segmented communication widens bid–ask spreads relative to no communication, while under unbiased information, full communication reduces trading volume by enabling faster agreement on fundamentals.

These findings carry implications for how communication environments shape financial markets. The consequences of expanding trader communication — through social media, chat rooms, or algorithmic information sharing — depend on the biases of the signals being transmitted. In fully connected networks, unbiased signals aggregate efficiently, while biased signals impede rather than promote information aggregation. Institutional responses aimed at improving market quality should therefore focus on the belief-formation process and, where information biases are suspected, consider whether restricting the reach of communication may be preferable to expanding it.

References

Acemoglu, D., Ozdaglar, A. and Siderius, J. (2024). A model of online misinformation. *Review of Economic Studies*, **91** (6), 3117–3150.

Charness, G., Oprea, R. and Yuksel, S. (2021). How do people choose between biased information sources? evidence from a laboratory experiment. *Journal of the European Economic Association*, **19** (3), 1656–1691.

Chen, H., De, P., Hu, Y. J. and Hwang, B.-H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *The Review of Financial Studies*, **27** (5), 1367–1403.

Choi, S., Galeotti, A. and Goyal, S. (2017). Trading in networks: theory and experiments. *Journal of the European Economic Association*, **15** (4), 784–817.

Cookson, J. A., Mullins, W. and Niessner, M. (2024). Social media and finance.

Corgnet, B., Deck, C., DeSantis, M., Hampton, K. and Kimbrough, E. O. (2022). When do security markets aggregate dispersed information? *Management Science*.

—, DeSantis, M. and Porter, D. (2024). Let's chat… when communication promotes efficiency in experimental asset markets. *Management Science*, **70** (10), 6550–6568.

Forsythe, R., Palfrey, T. R. and Plott, C. R. (1982). Asset valuation in an experimental market. *Econometrica: Journal of the Econometric Society*, pp. 537–567.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, **19** (4), 25–42.

Friedman, D., Harrison, G. W. and Salmon, J. W. (1984). The informational efficiency of experimental asset markets. *Journal of Political Economy*, **92** (3), 349–408.

Gentzkow, M., Shapiro, J. M. and Stone, D. F. (2015). Media bias in the marketplace: Theory. In *Handbook of media economics*, vol. 1, Elsevier, pp. 623–645.

Goldstein, I., Xiong, Y. and Yang, L. (2025). Information sharing in financial markets. *Journal of Financial Economics*, **163**, 103967.

Halim, E., Riyanto, Y. E. and Roy, N. (2019). Costly information acquisition, social networks, and asset prices: Experimental evidence. *The Journal of Finance*, **74** (4), 1975–2010.

Han, B. and Yang, L. (2013). Social networks, information acquisition, and asset prices. *Management Science*, **59** (6), 1444–1457.

Hirshleifer, D. (2020). Presidential address: Social transmission bias in economics and finance. *The Journal of Finance*, **75** (4), 1779–1831.

—, Peng, L. and Wang, Q. (2024). News diffusion in social networks and stock market reactions. *The Review of Financial Studies*, **38** (3), 883–937.

Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American economic review*, **92** (5), 1644–1655.

Jiao, P., Veiga, A. and Walther, A. (2020). Social media, news media and the stock market. *Journal of Economic Behavior & Organization*, **176**, 63–90.

Keasey, K., Lambrinoudakis, C., Mascia, D. V. and Zhang, Z. (2025). The impact of social media influencers on the financial market performance of firms. *European Financial Management*, **31** (2), 745–785.

Kuchler, T. and Stroebel, J. (2021). Social finance. *Annual Review of Financial Economics*, **13** (1), 37–55.

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

— and — (2021). How much information is incorporated into financial asset prices? experimental evidence. *The Review of Financial Studies*, **34** (9), 4412–4449.

Pedersen, L. H. (2022). Game on: Social networks and markets. *Journal of Financial Economics*, **146** (3), 1097–1119.

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

— and — (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica: Journal of the Econometric Society*, pp. 1085–1118.

Sunder, S. *et al.* (1992). *Experimental asset markets: A survey*. Carnegie Mellon University.

Walden, J. (2019). Trading, profits, and volatility in a dynamic information network model. *The Review of Economic Studies*, **86** (5), 2248–2283.

8 Appendix

Appendix A: Instructions

Here we present the instructions for the Biased Information in the Segmented Communication treatment. The instructions for the rest of the treatments can be found in the Online Appendix.

A1: Segmented Communication and Biased Information Treatment

Experiemntal instructions Welcome to the experiment. Please pay close attention and take your decisions carefully. By participating in this experiment, you can earn money. We will use an experimental currency unit that we call ECU, to measure your earnings. You will receive a fixed payment of 5 euros for participating in the experiment, plus a variable amount that depends on your performance in a market that we will describe in these instructions and on your responses to two sets of questions at the end of the experiment. At the end of the experiment, the amount of ECUs you have earned in one of the 12 rounds of the experiment will be randomly selected and converted into Euros. The conversion rate is 1,000 ECUs to 3 euros, and this will determine the variable part of your earnings. Therefore, your earnings in Euros will be higher, the higher your earnings in ECUs. At the end of the session, this amount of money will be paid to you privately in cash in an envelope.

If you have any questions or doubts, please raise your hand and we will come to answer you in private.

Any other type of communication with other participants is not allowed.

We will maintain your anonymity. You will be identified only by your identification number in our data. All information we collect will be used confidentially and solely for the purpose of this study

In this experiment, you will participate in a financial market in which you can buy and sell units of an asset. You will be one of 8 participants in the virtual financial market.

The experiment consists of 12 rounds, which are independent of one another. Each round of the experiment has two parts.

Part I: You will receive information about a financial asset and will be able to communicate with other participants via a chat. With this, you will able to form an opinion about the nature of the financial asset, which will be important in Part II of the experiment.

Part II: You will participate in a financial market where you can buy and sell units of an asset. In each round, every trader will have an initial endowment of 2,000 ECUs and 4 units of an asset. The ECUs that you earn in this part depend on your decisions, the decisions of other participants, and the outcome of a random draw of a ball.

At the end of these two parts, Part III of the experiment will begin. We will ask you to complete some questionnaires. The instructions for this part of the experiment will be explained on screens after Part II of the experiment.

Before the experiment begins, there will be 2 practice rounds so that you can familiarize yourself with the experiment before the 12 rounds of the main experiment. Note that the practice rounds will not be selected as payout rounds, but apart from this they are identical to the main rounds (the same rules and the same participants). Next, we explain the details of the main rounds and the other aspects of the experiment.

Information on the financial asset and communication

Neighbors

You will be connected with three other participants in your group (and they will also be connected with you and with each other). These participants will be called neighbors, and neighbors can communicate with each other. The neighbors will remain the same throughout the experiment. The other four participants will be neighbors among themselves.

Information

This part provides you with information about the financial asset. At the end of each round, the financial asset will have the same value for all participants. At the beginning of each round, none of the traders will know the value of the asset. The value of the asset will be determined in each round by the computer randomly choosing the color of a ball, which can be orange or green.

The color of the ball (orange or green) will be randomly selected and will determine the value of one unit of the asset at the end of one round:

- 50 ECUs if the selected ball is orange.
- 500 ECUs if the selected ball is green.

In each round, the color of the ball and, therefore, the value of the asset are identical for all participants in your group.

At the beginning of each round, you will be informed which of the two urns the computer will draw the ball from in that round. The urn will be the same for all participants in the experiment. In each round, we will alternate the selected urn so that in half of the rounds you have Urn 1 and in the other half Urn 2. The urns can have one of the following two compositions:

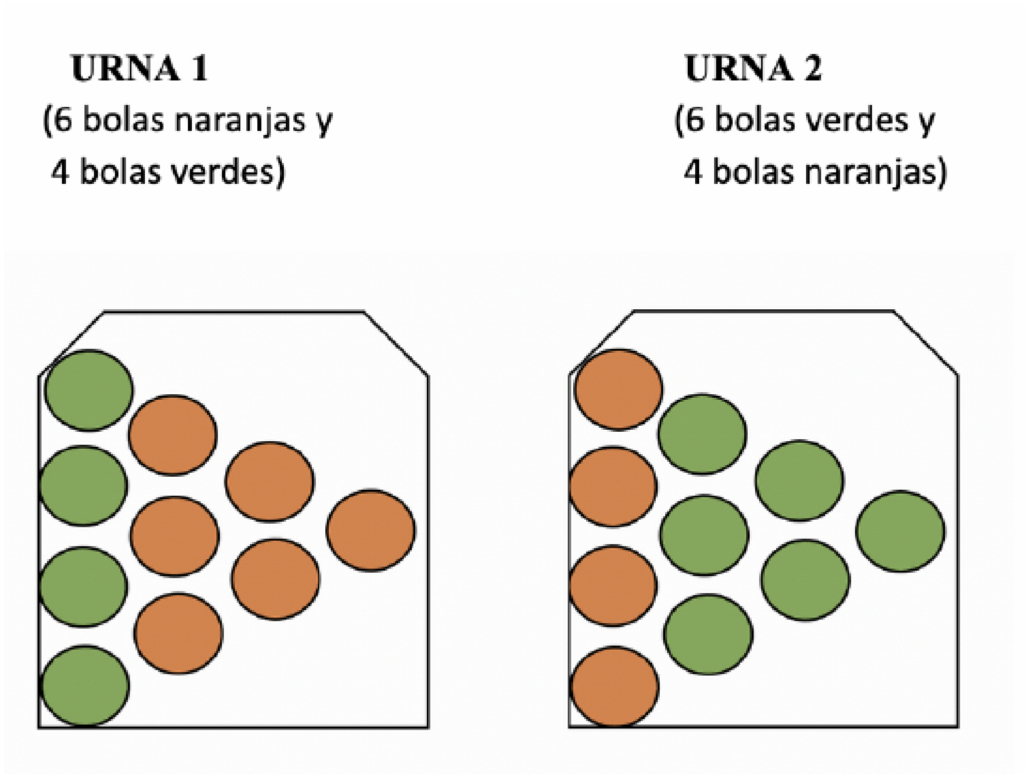

Both urns contain 10 balls, but their composition is different. Urn 1 has 6 orange balls and 4 green balls, while urn 2 has 6 green balls and 4 orange balls. With no additional information, you know that know that the ball is more likely to be orange if urn 1 has been selected and green if urn 2 has been chosen.

In each round, the computer will randomly draw a ball from the previously selected urn. All 10 balls have the same probability of being selected. In each round, the selected ball may be different and is the same for all participants in your group. Your objective is to infer whether the selected ball is orange or green, since this determines the value that the asset will pay at the end of each round of Part II. Given this, in Part II you will have the opportunity to buy and sell units of the asset at the prices proposed in the market. You will not know the color of the selected ball, but an advisor will provide you with information.

To help you discover the color of the ball, you will be given the opinion of a computerized advisor who knows the color of the selected ball. The advisor may say: “The ball is orange” or “The ball is green.” You and your neighbors will be given the opinion of one advisor (A or B), and the other participants (who are neighbors among themselves) will be given the opinion of the other advisor (A or B). Each participant will be assigned the same advisor in all rounds. At the beginning of each round, you will be reminded on your computer screen of the advisor who gives you information. The graph below shows the frequency of the opinion of the two advisors depending on the color of the ball that has been randomly taken out and determines the value of a unit of the asset. On the left is the graph for Advisor A and on the right is the graph for Advisor B.

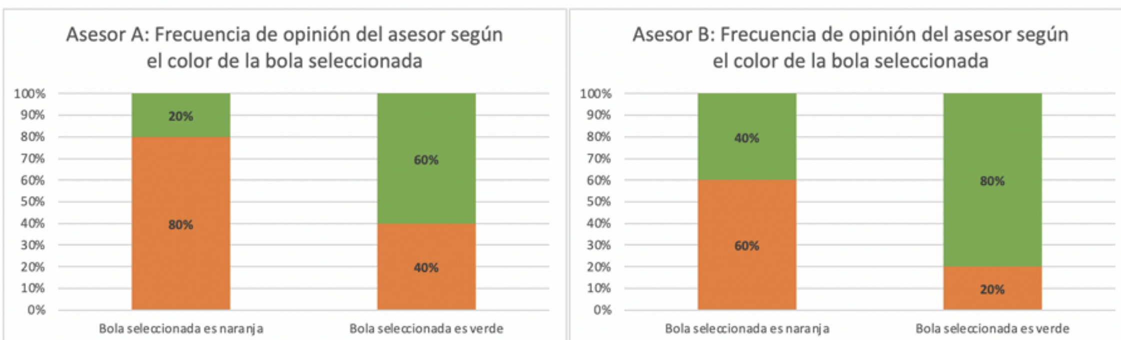

We can see that advisor A says the following:

- If the selected ball is orange:
  - Advisor A says: “The ball is orange” in 8 out of 10 cases.
  - Advisor A says: “The ball is green” in 2 out of 10 cases.
- If the selected ball is green:
  - Advisor A says: “The ball is green” in 6 out of 10 cases.
  - Advisor A says: “The ball is orange” in 4 out of 10 cases.

However, advisor B says the following:

- If the selected ball is orange:
  - Advisor B says: "The ball is orange" in 6 out of 10 cases.
  - Advisor B says: "The ball is green" in 4 out of 10 cases.
- If the selected ball is green:
  - Advisor B says: "The ball is green" in 8 out of 10 cases.
  - Advisor B says: "The ball is orange" in 2 out of 10 cases.

Important: The opinion given by the same advisor on the color of the ball may differ from participant to participant.

To help you with the task of figuring out the color of the ball, we then provide some tables in which we calculate the probability that the ball is green or orange depending on the urn assigned to you in each round and the advisor. For advisor A:

Asesor A	Urna 1 6 bolas naranjas; 4 bolas verdes	Urna 2 4 bolas naranjas; 6 bolas verdes
El asesor dice: "La bola es naranja"		
Probabilidad que la bola sea naranja	75%	57%
Probabilidad que la bola sea verde	25%	43%
El asesor dice: "La bola es verde"		
Probabilidad que la bola sea naranja	*33%*	18%
Probabilidad que la bola sea verde	67%	82%

For advisor B:

Asesor B	Urna 1 6 bolas naranjas; 4 bolas verdes	Urna 2 4 bolas naranjas; 6 bolas verdes
El asesor dice: "La bola es naranja"		
Probabilidad que la bola sea naranja	82%	67%
Probabilidad que la bola sea verde	18%	33%
El asesor dice: "La bola es verde"		
Probabilidad que la bola sea naranja	43%	25%
Probabilidad que la bola sea verde	57%	75%

In each round, we will display these probabilities on the screen given the selected urn and the advisor's message.

Communication

Before trading, you will have the opportunity to communicate with your neighbors. Remember that you have 3 neighbors. You and your neighbors all have the same advisor. The other four participants will be neighbors, will be able to communicate with one another, and will have the other advisor.

You will be able to communicate with your neighbors via a chat for 2 minutes in each round. You can write whatever you want in the chat, as long as you are respectful of others and do not reveal your identity to preserve the anonymity of the experiment. On the screen you will be able to see a clock with remaining chat time.

Example of the screens

Below, we provide an example of the first screen of the experiment. The first screen of each round will inform you of the advisor (who is the same for all rounds) and of the the urn selected in the round (which changes from round to round). You do not have to make any decision on this screen. The advisor's message and the probability that the ball is orange or green will also appear. Simply click "Next" once you have read the information.

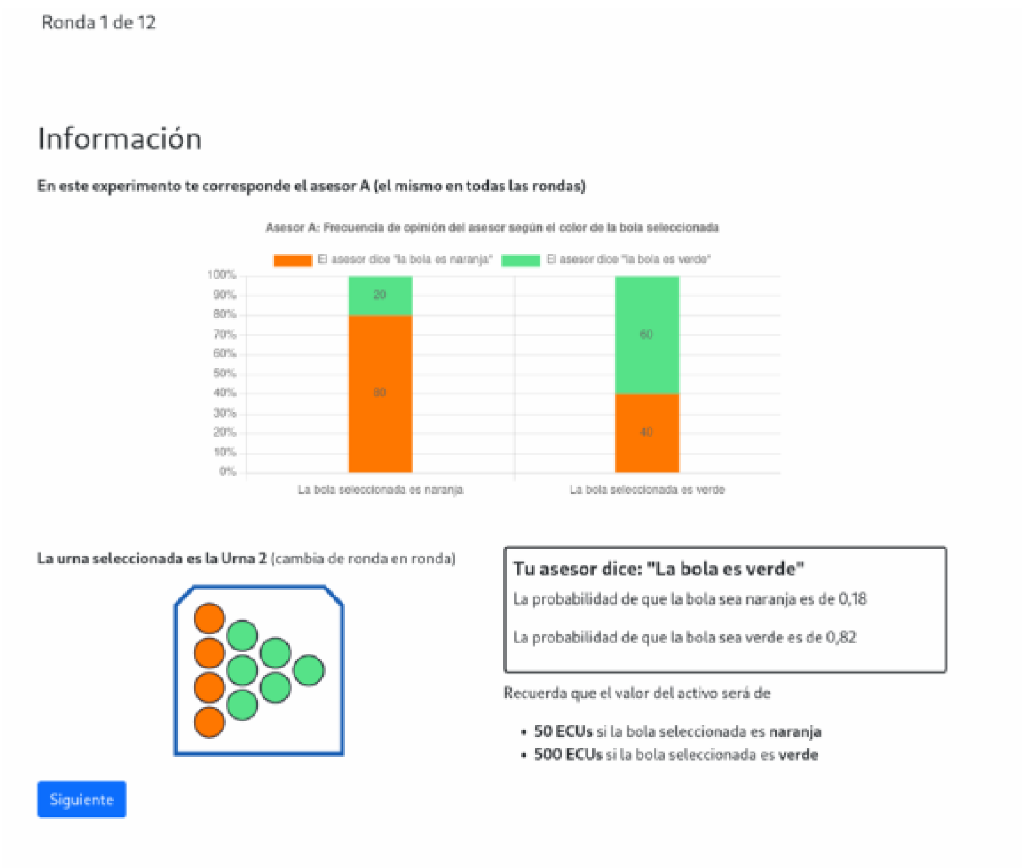

The second screen of each round is the chat screen, where you will be able to communicate with your neighbors and your neighbors will be able to communicate with you. Here we also remind you of the number of neighbors (that is, the participants you are connected with). You will see that at the end of the communication screen, and after communicating with your neighbors, you will be asked about the probability with which you estimate that the ball chosen was green or orange.

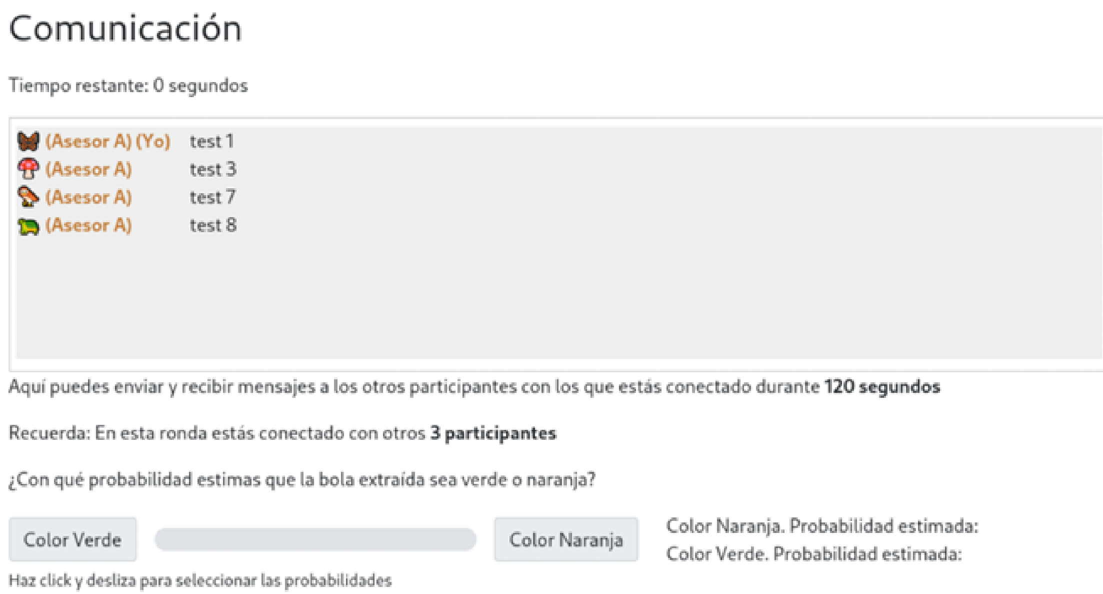

Part II: Trading in the market

In Part II of the experiment, you will be able to buy and sell units of the asset with other participants in your group of 8 (including yourself). The group will remain the same throughout the entire experiment. Remember that all traders have the same endowment per round, consisting of 2,000 ECUs and 4 units of the asset. Also remember that at the end of the round, the value of one unit of the asset will be:

- 50 ECUs if the chosen ball is orange.
- 500 ECUs if the chosen ball is green.

In each round, your earnings come from two sources: (i) the value of the units of the asset you hold at the end of the round, and (ii) the ECUs you obtain from buying and selling units of the asset. The units of the asset and the ECUs you hold at the end of a round cannot be used in the next round. In the experiment there are 12 rounds independent of each another, so the selected urn (and therefore the value of the asset) may differ across rounds. During each trading round, you can buy or sell as many units of the asset as you wish, given that you follow the rules specified below.

To buy an asset, you have the following options:

1. You can accept a sell offer from other traders listed in the list of sell offers. Note that a sell offer always consists of one unit of the asset being available for purchase and specifies the price at which you can buy that unit. For example, if you accept another trader's sell offer at 120 ECUs, your ECU balance will decrease by 120 and you will hold one additional unit of the asset or good.

2. You can create a buy offer that specifies a price at which other traders can sell one unit of the asset to you. Note that you can only buy a unit of the asset if your buy offer is accepted by the other trader. If no one accepts your buy offer, you do not obtain additional units of the asset. Also note that if you wish to buy more than one unit of the asset, you would have to accept multiple sell offers or create multiple buy offers.

Similar to buying a unit of asset, to sell a unit of asset:

1. You can accept a buy offer from other traders listed in the list of buy offers. For example, if you accept a buy offer at a price of 250 ECUs, your ECU balance will increase by 250 and you will hold one fewer unit of the asset. Similarly, each buy offers refers to only one unit of the asset or good.

2. You can create a sell offer that specifies a price at which other traders can buy one unit of the asset from you. Note that you can only sell a unit of the asset if your sell offer is accepted. If no one accepts your sell offer, the number of units of the asset you hold does not decrease.

Similarly if you want to sell more than one unit of the asset you will have to accept more offers of purchase or create more offers of sale. At all times you cannot spend more ECU or sell more units of the asset than you have at that time.

Throughout trading, you can see all active offers and corresponding prices. To create an offer simply press the "Add offer" button in the corresponding column and then a window will appear where you can enter a buy or sell price, respectively.

In general, you must accept the highest purchase offers and the lowest sales offers. To facilitate viewing, the offer listings will be arranged in such a way that the purchase offers appear from high to low and sale offers from low to high (i.e. best deals will appear above).

The following screenshot shows an example of the experiment's third screen, the trading screen:

Your own buy and sell offers will appear in the list of offers in blue rather than in black. You cannot accept your own offers, but you can select them and click "delete" if you wish. Offers disappear from the list of offers when someone accepts them or withdraws them.

If there is any buy or sell offer that you can no longer make (either because you do not have units of the asset for your sell offer or because you do not have enough ECUs for your buy offer), it will also disappear from the list.

After you complete a purchase or sale, your ECU balance and the number of assets you hold will be updated accordingly.

In each round, you will have 3 minutes to trade. The remaining time will be displayed (in seconds) at the top right of the screen.

The final screen concludes the round. It will display the average transaction price during the round and your earnings. Your earnings for each round are the sum of the ECUs you have at the end of the round and the value of the units of the asset you hold at the end of the round, in ECUs. A new round will begin once all traders have clicked "Next" on the screen.

Start of the Experiment

If you have questions that cannot be fully answered by the instructions, please raise your hand and ask for help before proceeding. You can consult these instructions at all times, but please do not talk to other participants. Good luck with the experiment, and thank you for your participation.

Part III: Questionnaires

There will be some questionnaires that we will ask you to respond. The specific instructions will be on the screen.

A2: Unbiased Information Treatment: Information Section

This reports the main changes to the information section for the Unbiased information treatments. For full instructions, see the Online Appendix.

Part I: Information on the financial asset

Information

This part provides you with information about the financial asset. At the end of each round, the financial asset will have the same value for all participants. At the beginning of each round, none of the traders will know the value of the asset. The value of the asset will be determined in each round by the computer randomly choosing the color of a ball, which can be orange or green.

The color of the ball (orange or green) will be randomly selected and will determine the value of one unit of the asset at the end of one round:

- 50 ECUs if the selected ball is orange.
- 500 ECUs if the selected ball is green.

In each round, the color of the ball and, therefore, the value of the asset are identical for all participants in your group.

At the beginning of each round, you will be informed which of the two urns the computer will draw the ball from in that round. The urn will be the same for all participants in the experiment. In each round, we will alternate the selected urn so that in half of the rounds you have Urn 1 and in the other half Urn 2. The urns can have one of the following two compositions:

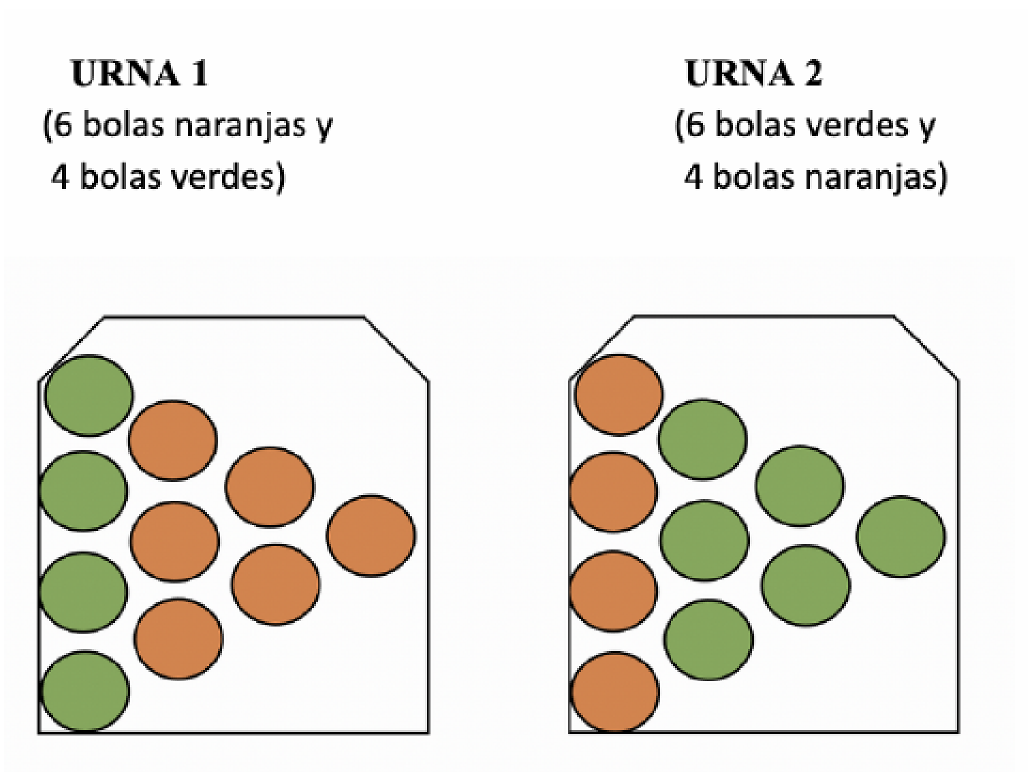

Both urns contain 10 balls, but their composition is different. Urn 1 has 6 orange balls and 4 green balls, while urn 2 has 6 green balls and 4 orange balls. With no additional information, you know that know that the ball is more likely to be orange if urn 1 has been selected and green if urn 2 has been chosen.

In each round, the computer will randomly draw a ball from the previously selected urn. All 10 balls have the same probability of being selected. In each round, the selected ball may be different and is the same for all participants in your group. Your objective is to infer whether the selected ball is orange or green, since this determines the value that the asset will pay at the end of each round of Part II. Given this, in Part II you will have the opportunity to buy and sell units of the asset at the prices proposed in the market. You will not know the color of the selected ball, but an advisor will provide you with information.

To help you discover the color of the ball, you will be given the opinion of a computerized advisor who knows the color of the selected ball. The advisor may say: "The ball is orange" or "The ball is green." You and three other participants will be given the opinion of one advisor (A or B), and the other four participants will be given the opinion of the other advisor (B or A). Each participant will be assigned the same advisor in all rounds. At the beginning of each round, you will be reminded on your computer screen of the advisor who gives you information. The graph below shows the frequency of the opinion of the two advisors depending on the color of the ball that has been randomly taken out and determines the value of a unit of the asset. On the left is the graph for Advisor A and on the right is the graph for Advisor B.

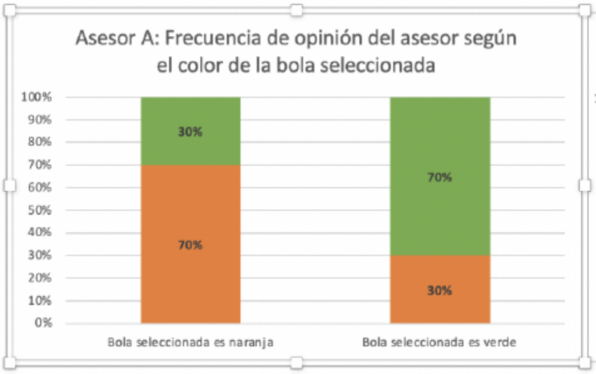

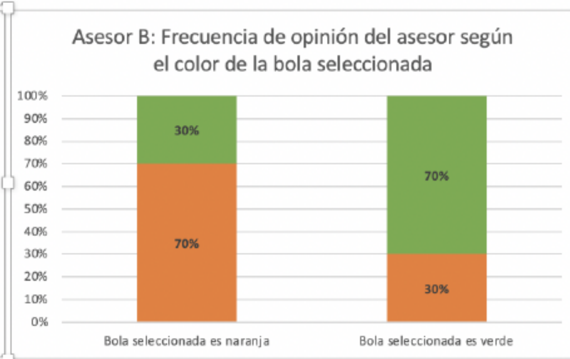

We can see that advisor A says the following:

- If the selected ball is orange:
  - Advisor A says: "The ball is orange" in 7 out of 10 cases.
  - Advisor A says: "The ball is green" in 3 out of 10 cases.
- If the selected ball is green:
  - Advisor A says: "The ball is green" in 7 out of 10 cases.
  - Advisor A says: "The ball is orange" in 3 out of 10 cases.

However, advisor B says the following:

- If the selected ball is orange:
  - Advisor B says: "The ball is orange" in 7 out of 10 cases.
  - Advisor B says: "The ball is green" in 3 out of 10 cases.
- If the selected ball is green:
  - Advisor B says: "The ball is green" in 7 out of 10 cases.
  - Advisor B says: "The ball is orange" in 3 out of 10 cases.

Important: The opinion given by the same advisor on the color of the ball may differ from participant to participant.

To help you with the task of figuring out the color of the ball, we then provide some tables in which we calculate the probability that the ball is green or orange depending on the urn assigned to you in each round and the advisor. For advisor A:

Asesor A	Urna 1 6 bolas naranjas; 4 bolas verdes	Urna 2 4 bolas naranjas; 6 bolas verdes
El asesor dice: "La bola es naranja"		
Probabilidad que la bola sea naranja	78%	61%
Probabilidad que la bola sea verde	22%	39%
El asesor dice: "La bola es verde"		
Probabilidad que la bola sea naranja	39%	22%
Probabilidad que la bola sea verde	61%	78%

For advisor B:

Asesor B	Urna 1 6 bolas naranjas; 4 bolas verdes	Urna 2 4 bolas naranjas; 6 bolas verdes
El asesor dice: "La bola es naranja"		
Probabilidad que la bola sea naranja	78%	61%
Probabilidad que la bola sea verde	22%	39%
El asesor dice: "La bola es verde"		
Probabilidad que la bola sea naranja	39%	22%
Probabilidad que la bola sea verde	61%	78%

To help you, in each round, we will display these probabilities on the screen given the selected urn and the advisor's message.

Example of the screens

Below, we provide an example of the first screen of the experiment. The first screen of each round will inform you of the advisor (who is the same for all rounds) and of the the urn selected in the round (which changes from round to round). You do not have to make any decision on this screen. The advisor's message and the probability that the ball is orange or green will also appear. Simply click "Next" once you have read the information.

On the second screen of each round, we will ask you for the probability with which you estimate that the selected ball was green or orange.

Appendix B: Chat analysis

B.1. Coder's Protocol for Analysing Chat Data

Analysis of the chat for conversation (treatment, round, group/sub-group)

1. General

   i. Majorly pro-social (Hello, How are you) Yes – 1 or No – 0.

   [Prosocial_coder]*

   ii. Majorly anti-social (any other insult or accusation) Yes – 1 or No – 0.

   [Anti-social_coder]

   iii. Boredom. Yes – 1 or No – 0.

   [Boredom_coder]

2. How many people in the group (4 or 8) say something in the conversation?

   [Chat_Num_Participants_coder]

3. Do they share information about the signal (advisor) or your belief about the color of the ball? Yes – 1 or No — 0.

   [Info_sharing_ind_coder]

[13]* In the analysis **coder** is represented as 1 (Human coder), 2 (Human coder) and ChatGPT as the coder.

a. How many (out of 4 or 8) share information about the signal or their belief?

[Num_share_info_ind_coder]

b. If you answered yes to question 3: Do they share the information about what the advisor said? Yes (1) or No (0).

[Share_advisor_info_ind_coder]

i. Is there any other message with false information? Yes (1), or No (0).

[False_info_ind_coder]

ii. How many (out of 4 or 8) provide false information?

[Num_share_false_info_ind_coder]

c. If you have answered yes to question 3. Do they share information on their beliefs on the color of the ball ? Yes (1) or No (0).

[Share_beliefs_ind_coder]

d. If you have answered Yes to 3, and No to b and c). It cannot be distinguished whether they are talking about the advisor's information or about their beliefs (mainly). Yes (1) or No (0).

[Share_difficulty_interpret_ind_coder]

4. Do they share information on market strategies? Yes – 1 or No – 0.

[Share_trading_strategies_coder]

i. Do they share information on what to do if the ball is green? Yes (1) or No (0).

[Share_trading_green_coder]

ii. Do they share information on what to do if the ball is orange? Yes (1) or No (0).

[Share_trading_orange_coder]

iii. Do they share intentions about market trading? Yes (1) or No (0).

[Share_intentions_coder]

iv. Do they ask others questions about their market strategies? Yes (1) or No (0).

[Ask_trading_strategies_coder]

v. Others (add a note on the type of strategy)

5. Expression of confusion or clarifications, doubts, or questions about the instructions/rules of the game?

[Confusion_game_coder]

i. On the color of the ball Yes – 1 or No – 0

[Doubts_color_ball_coder]

ii. On the urn. Yes – 1 or No – 0

[Doubts_urn_coder]

iii. On what your advisor said. Yes – 1 or No – 0

[Doubts_own_advisor_coder]

iv. On what both your advisors said. Yes – 1 or No – 0

[Doubts_other_advisor_coder]

v. In the 2-Island treatments (group chat of 4), on what happened in the other island. Yes – 1 or No – 0

[Doubts_other_island_coder]

vi. The rules of the game. Yes – 1 or No – 0

vii. The better strategy to play the game. Yes – 1 or No – 0

[Doubts_rulesofgame_coder]

6. Do they make a reference with respect to learning over time? Yes – 1 or No – 0.

   [Learning_frompast_coder]

7. Do they refer to their past earnings? Yes – 1 or No – 0.

   [Past_earnings_coder]

8. Do they infer anything from possible past strategies that can influence present or future strategies? Yes – 1 or No – 0.

   [Past_strategies_coder]

9. Was a consensus (majority) reached about the color of the ball? Yes or No.

   [Consensus_colorball_coder]

   a. Is it true? Yes or No.

      [True_Consensus_coder]

   b. What percentage of the group (4 or 8) reaches this majority/consensus?

      [Percent_Consensus_coder]

10. Was a consensus (majority) reached about the market trading strategy? Yes – 1 or No – 0.

    [Consensus_tradingstrategy_coder]

    a. Which one?

       [Consensus_tradingstrategy_which_coder]

11. Do they propose to coordinate or do they coordinate over time? Yes – 1 or No – 0.

    [Coordination_coder]

    a. Which one?

12. Other comments.

B.2. Coders' agreement analysis

Table 12
Cohen's Kappa Statistics: Human Coder 1 and 2

Variable	Agree (%)	Exp. Agree (%)	κ	Std. Err.	Z	p-value
Num. Share Info	75.69	25.79	0.6725	0.0223	30.22	<0.001
Info Sharing	98.15	92.44	0.7552	0.0478	15.82	<0.001
False Info	93.29	50.50	0.8644	0.0481	17.97	<0.001
Num. Share False Info	84.26	37.31	0.7489	0.0302	24.80	<0.001
Share Trading Strat.	87.04	57.50	0.6950	0.0475	14.64	<0.001
Consensus Colorball	82.18	50.41	0.6406	0.0481	13.32	<0.001
True Consensus	86.11	51.14	0.7157	0.0481	14.89	<0.001

Note: All kappa statistics significant at $p < 0.001$.

Table 13
Cohen's Kappa Statistics: Human Coder 1 and 2

Variable	Cohen's κ	Std. Error	p-value
Num. Share Info	0.672	0.022	<0.001
False Info	0.864	0.048	<0.001
Num. Share False Info	0.749	0.030	<0.001
Share Trading Strat.	0.695	0.047	<0.001
Consensus Colorball	0.641	0.048	<0.001
True Consensus	0.716	0.048	<0.001
Info Sharing	0.755	0.048	<0.001
Chat Num. Participants	0.983	0.034	<0.001
Share Intentions	0.585	0.046	<0.001
Doubts Urn	0.708	0.048	<0.001
Doubts Own Advisor	0.549	0.047	<0.001
Past Earnings	0.709	0.047	<0.001
Consensus Trading Strat.	0.564	0.048	<0.001
Observations	432		

Note: Cohen's kappa statistics reported with standard errors and p-values.

9 Appendix C: Further Analysis

9.1 Distribution of transaction prices across treatments

Table 14
Distribution of Trades Across Price Bins and Treatments

pricebin	*B_AllCom*	*B_NoCom*	*B_SegCom*	*U_AllCom*	*U_NoCom*	*U_SegCom*	*Total*
0–<50	8	5	51	10	32	62	168
	(0.60)	(0.36)	(3.60)	(0.96)	(1.71)	(4.05)	(1.96)
50–<=500	1,302	1,324	1,308	1,007	1,813	1,466	8,220
	(97.97)	(96.57)	(92.31)	(97.11)	(96.95)	(95.75)	(96.08)
500–<=600	5	17	46	11	7	1	87
	(0.38)	(1.24)	(3.25)	(1.06)	(0.37)	(0.07)	(1.02)
600–<=700	4	5	5	5	15	1	35
	(0.30)	(0.36)	(0.35)	(0.48)	(0.80)	(0.07)	(0.41)
>700	10	20	7	4	3	1	45
	(0.75)	(1.46)	(0.49)	(0.39)	(0.16)	(0.07)	(0.53)
Total	1,329	1,371	1,417	1,037	1,870	1,531	8,555
	(100.00)	(100.00)	(100.00)	(100.00)	(100.00)	(100.00)	(100.00)

Note. Percentages in parentheses.

9.2 Transaction Prices in each treatment

Figure 5: Transaction Price B_AllCom

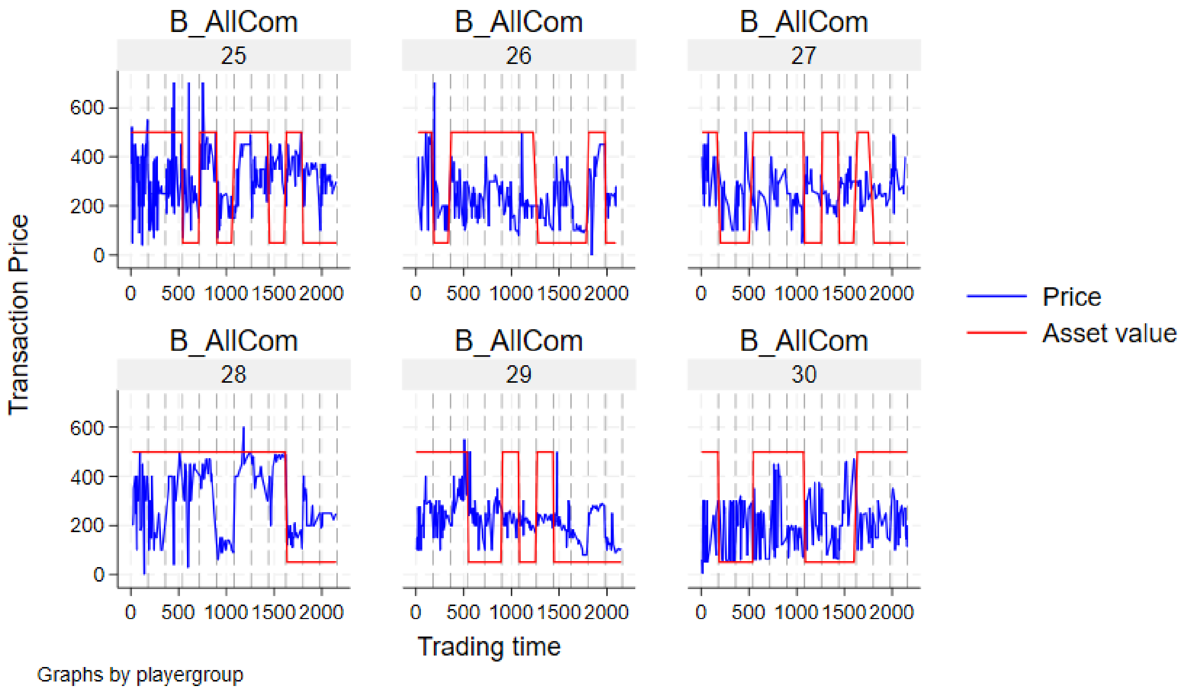

Figure 6: Transaction Price U_AllCom

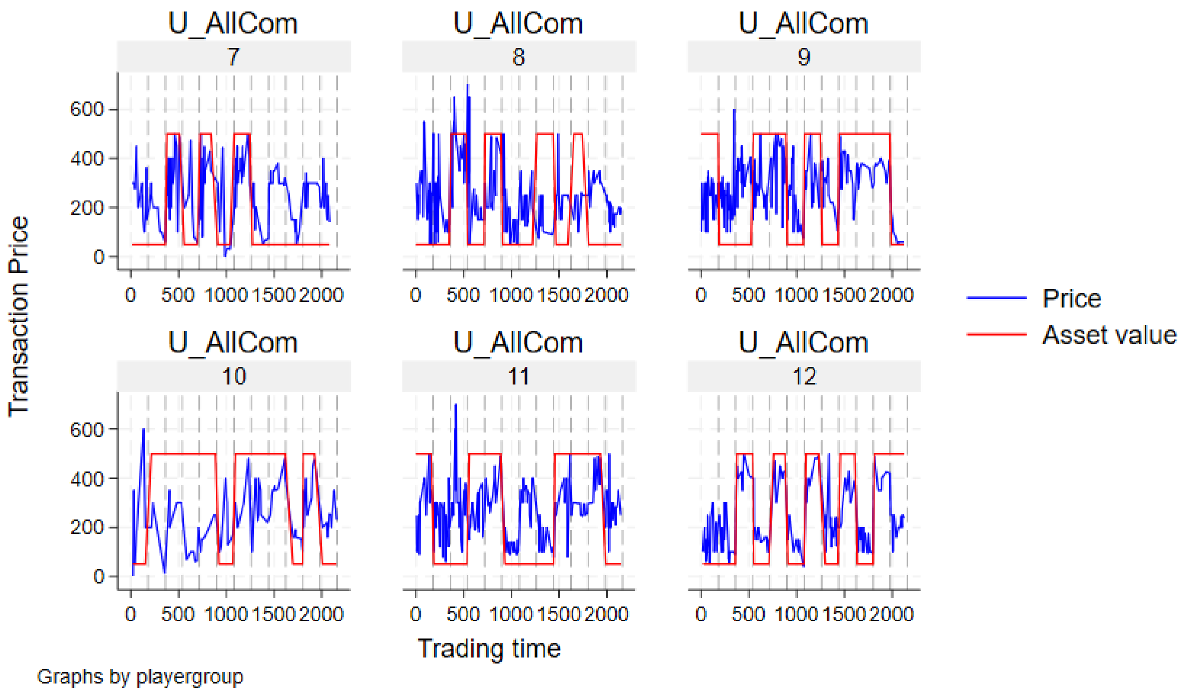

Figure 7: Transaction Price B_SegCom

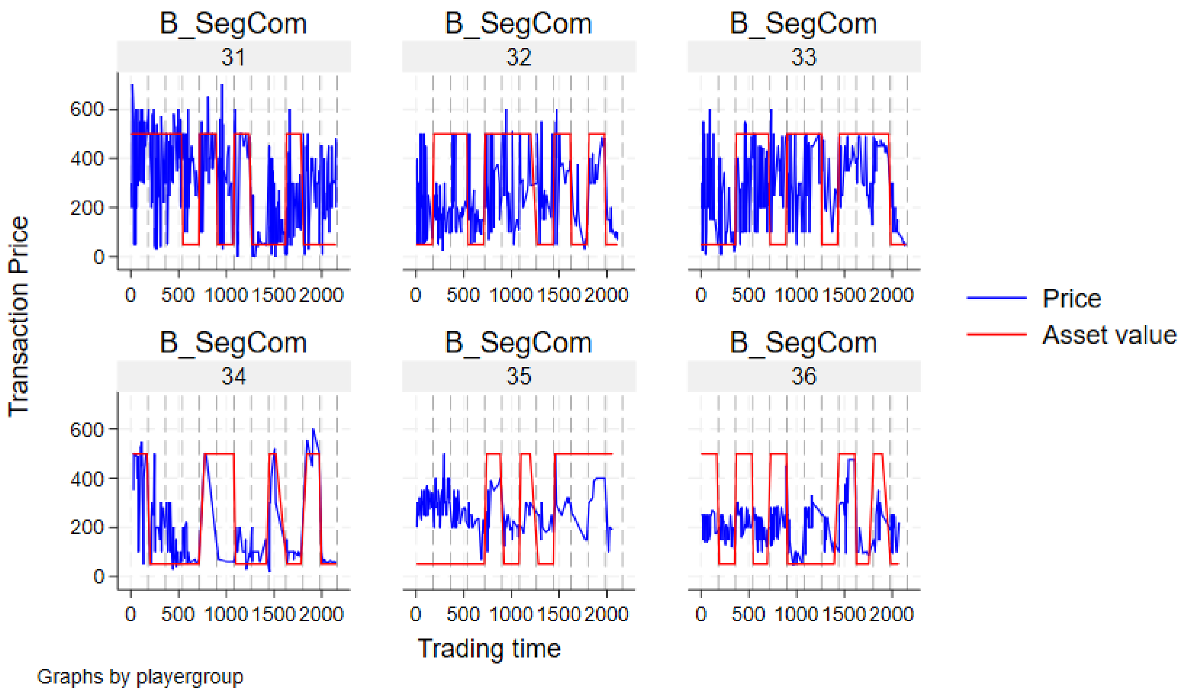

Figure 8: Transaction Price U_SegCom

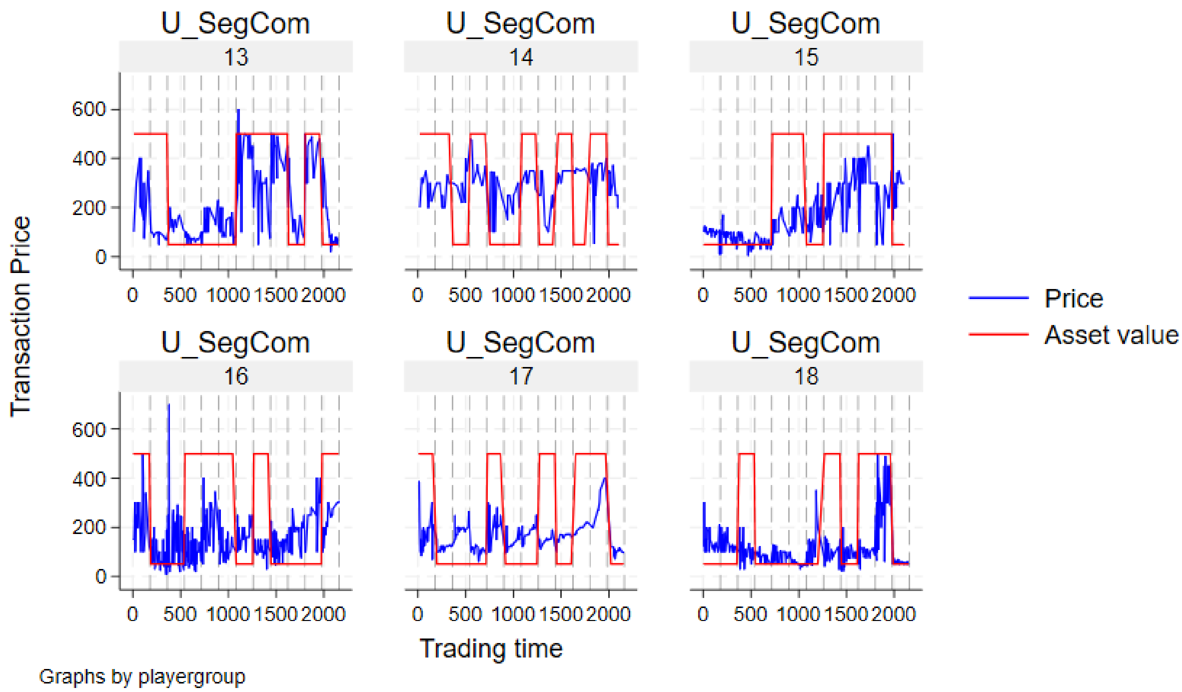

Figure 9: Transaction Price U_NoCom

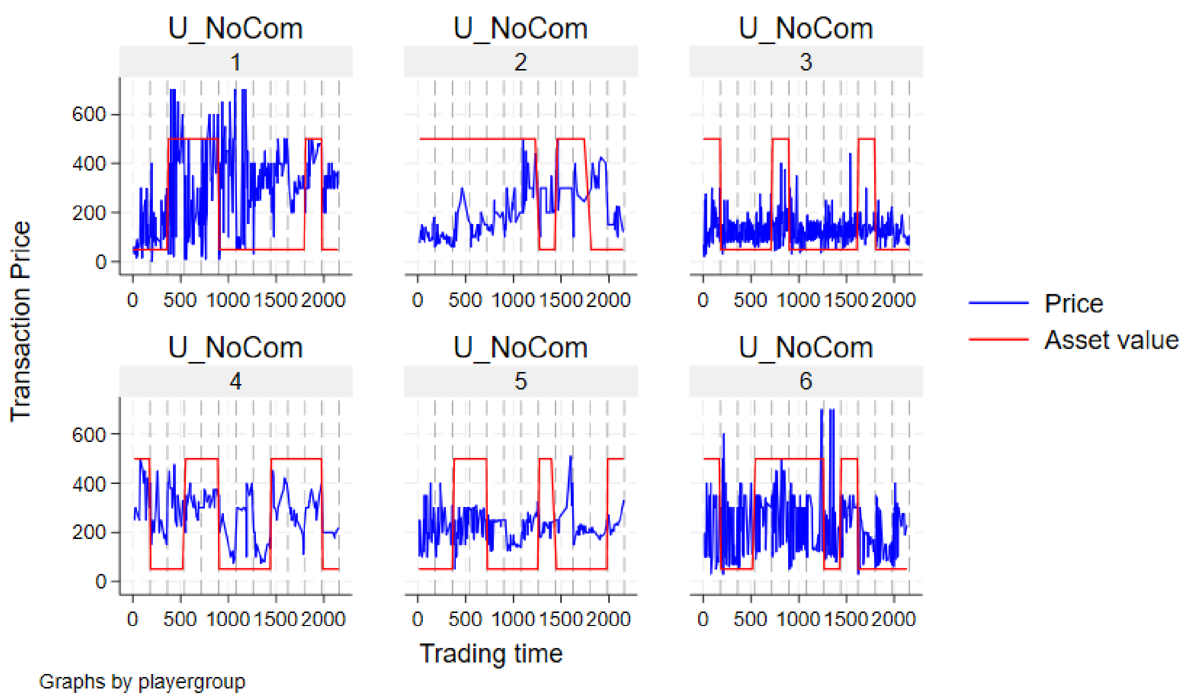

Figure 10: Transaction Price B_NoCom

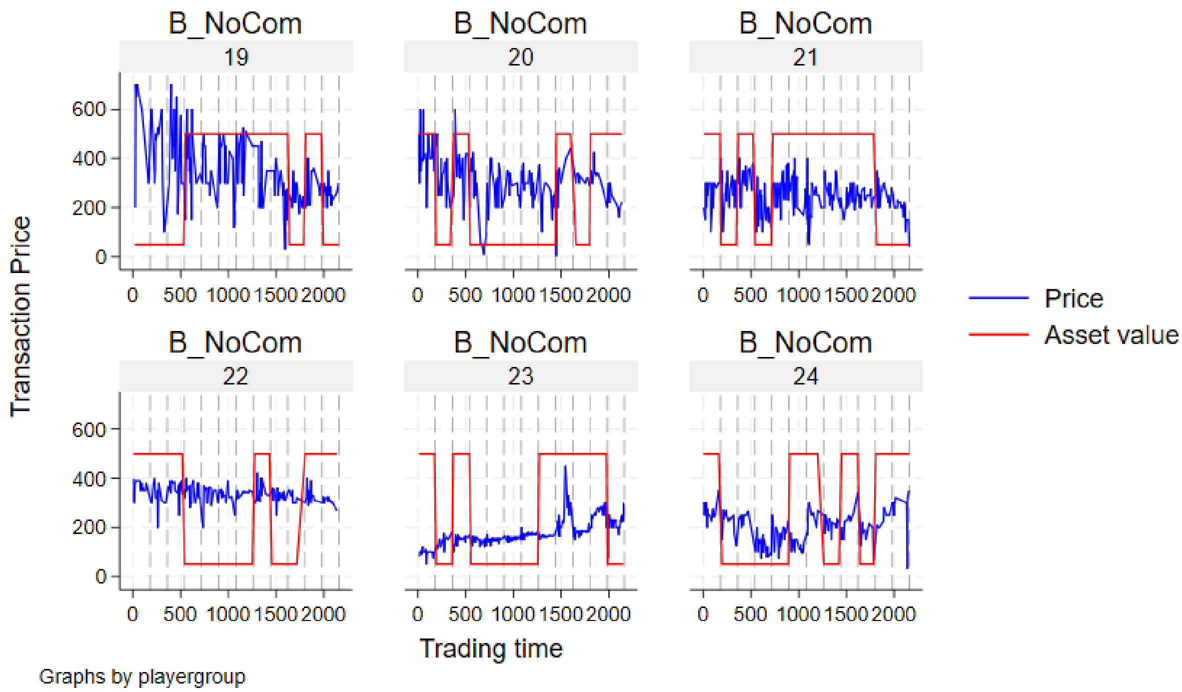

Table 15
Distribution of Actions Across Treatments

Action	*B_AllCom*	*B_NoCom*	*B_SegCom*	*U_AllCom*	*U_NoCom*	*U_SegCom*	**Total**
Trade	1,329	1,371	1,417	1,037	1,870	1,531	8,555
	39.72%	42.79%	39.46%	37.13%	48.41%	46.39%	42.57%
Withdraw	814	894	758	630	695	718	4,509
	24.33%	27.90%	21.11%	22.56%	17.99%	21.76%	22.44%
Submitted but not Executed	1,203	939	1,416	1,126	1,298	1,051	7,033
	35.95%	29.31%	39.43%	40.32%	33.60%	31.85%	35.00%
Total	3,346	3,204	3,591	2,793	3,863	3,300	20,097
	100%	100%	100%	100%	100%	100%	100%

Note. Percentages represent column percentage.

9.3 Distribution of Submitted Prices Across Treatments and Actions

Table 16
Distribution of Submitted Prices Across Price Bins across Treatments

pricebin	*B_AllCom*	*B_NoCom*	*B_SegCom*	*U_AllCom*	*U_NoCom*	*U_SegCom*	*Total*
0–<50	199	54	308	146	144	195	1,046
	(5.95)	(1.69)	(8.58)	(5.23)	(3.73)	(5.91)	(5.20)
50–<=500	2,906	2,831	2,972	2,373	3,490	2,948	17,520
	(86.85)	(88.36)	(82.76)	(84.96)	(90.34)	(89.33)	(87.18)
500–<=700	94	112	206	155	101	39	707
	(2.81)	(3.50)	(5.74)	(5.55)	(2.61)	(1.18)	(3.52)
700–<=1000	69	119	34	42	37	23	324
	(2.06)	(3.71)	(0.95)	(1.50)	(0.96)	(0.70)	(1.61)
1000–<=2000	17	39	16	13	13	5	103
	(0.51)	(1.22)	(0.45)	(0.47)	(0.34)	(0.15)	(0.51)
>2000	61	49	55	64	78	90	397
	(1.82)	(1.53)	(1.53)	(2.29)	(2.02)	(2.73)	(1.98)
Total	3,346	3,204	3,591	2,793	3,863	3,300	20,097
	(100.00)	(100.00)	(100.00)	(100.00)	(100.00)	(100.00)	(100.00)

Note. Percentages in parentheses represent column percentages within treatment.

Table 17
Communication Effects on Pre-Trade Belief Absolute Error (excluding price above 700)

	(1)	(2)	(3)	(4)
Biased × *SegCom*	−0.064***	−0.064***	−0.065***	−0.043***
	(0.02)	(0.02)	(0.02)	(0.02)
Biased × *AllCom*	0.011	0.011	0.013	0.010
	(0.02)	(0.02)	(0.02)	(0.02)
Unbiased × *NoCom*	0.017	0.017	0.014	0.011
	(0.01)	(0.01)	(0.01)	(0.01)
Unbiased × *SegCom*	−0.033	−0.033	−0.035	−0.038
	(0.02)	(0.02)	(0.02)	(0.02)
Unbiased × *AllCom*	−0.047**	−0.047**	−0.049**	−0.029*
	(0.02)	(0.02)	(0.02)	(0.02)
Round		0.000	−0.001	−0.001
		(0.00)	(0.00)	(0.00)
High Fundamental			−0.034***	−0.034***
			(0.01)	(0.01)
More Info. Signal				−0.339***
				(0.02)
Constant	0.431***	0.433***	0.453***	0.636***
	(0.01)	(0.02)	(0.02)	(0.02)
Observations	432	432	432	432
R^2 (overall)	0.044	0.045	0.059	0.396

Note. Standard errors clustered at the player-group level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

9.4 Further analysis on Beliefs

Table 18 reports pairwise contrasts for post-trade belief accuracy and belief dispersion. The overall interaction is highly significant for post-trade belief accuracy ($\chi^2(5) = 31.52$, $p < 0.0001$), mirroring the pre-trade belief pattern: under biased information, SegCom produces significantly more accurate post-trade beliefs than both NoCom (coef. = 0.086, $p_{\text{Holm}} = 0.002$) and AllCom (coef. = -0.096, $p_{\text{Holm}} = 0.004$), while under unbiased information, AllCom outperforms NoCom (coef. = 0.066, $p_{\text{Holm}} < 0.001$). By contrast, belief dispersion shows no significant variation across treatments ($\chi^2(5) = 5.17$, $p = 0.395$), suggesting that communication and information structures shape the accuracy of beliefs but not the degree of disagreement among traders within a group.

Table 18

Pairwise Comparisons of Post-Trade Belief Accuracy and Belief Dispersion across Information Structure and Communication Structure

	Coef.	SE	p (unadj.)	p (Holm)
Panel A: Post-Trade Belief Accuracy				
Network structure comparison within Biased Info				
NoCom vs. SegCom	0.0865**	0.0255	0.001	0.002
NoCom vs. AllCom	−0.0090	0.0235	0.701	0.701
SegCom vs. AllCom	−0.0955**	0.0308	0.002	0.004
Network structure comparison within Unbiased Info				
NoCom vs. SegCom	0.0465†	0.0217	0.032	0.064
NoCom vs. AllCom	0.0656**	0.0166	0.000	0.000
SegCom vs. AllCom	0.0191	0.0259	0.462	0.462
Information bias comparisons within each network structure				
NoCom: Biased vs. Unbiased	−0.0010	0.0129	0.940	0.940
SegCom: Biased vs. Unbiased	−0.0409	0.0309	0.185	0.371
AllCom: Biased vs. Unbiased	0.0737**	0.0260	0.005	0.013
Overall interaction: $\chi^2(5) = 31.52$, $p < 0.0001$				
Panel B: Belief Dispersion (SD)				
Network structure comparison within Biased Info				
NoCom vs. SegCom	0.0221	0.0258	0.390	0.780
NoCom vs. AllCom	−0.0231	0.0297	0.438	0.438
SegCom vs. AllCom	−0.0452†	0.0266	0.090	0.270
Network structure comparison within Unbiased Info				
NoCom vs. SegCom	0.0256	0.0191	0.180	0.541
NoCom vs. AllCom	0.0140	0.0191	0.466	0.932
SegCom vs. AllCom	−0.0117	0.0219	0.595	0.595
Information bias comparisons within each network structure				
NoCom: Biased vs. Unbiased	−0.0133	0.0233	0.569	1.000
SegCom: Biased vs. Unbiased	−0.0098	0.0220	0.657	0.657
AllCom: Biased vs. Unbiased	0.0237	0.0265	0.371	1.000
Overall interaction: $\chi^2(5) = 5.17$, $p = 0.395$				

Notes: All contrasts estimated from single random-effects GLS regressions ($N = 432$, 36 groups) with robust standard errors clustered at the group level. Panel A: dependent variable is the mean absolute deviation between post-trade belief and the fundamental value. Panel B: dependent variable is the within-group standard deviation of beliefs. Holm–Bonferroni correction applied separately within each hypothesis family. Marginal cell means — Panel A: *Biased Info*: NoCom = 0.438, SegCom = 0.351, AllCom = 0.447. *Unbiased Info*: NoCom = 0.439, SegCom = 0.393, AllCom = 0.373. Panel B: *Biased Info*: NoCom = 0.214, SegCom = 0.192, AllCom = 0.237. *Unbiased Info*: NoCom = 0.228, SegCom = 0.202, AllCom = 0.214. $^{**}p < 0.05$ after Holm correction; $^{\dagger}p < 0.10$ unadjusted only.