NBER WORKING PAPER SERIES

WAITING FOR THE RIGHT OFFER:
LABORATORY EVIDENCE ON HOW NEWS AFFECTS BARGAINING

Tingting Ding
Steven F. Lehrer

Working Paper 34543
http://www.nber.org/papers/w34543

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2025

The research reported in the paper was approved by the Research Ethics Board of NYU under IRB FY2016-1041. Ding initiated this project while at the Shanghai University of Finance and Economics and thanks the Natural Science Foundation of China (grant number 72173076) for financial support. Lehrer thanks SSHRC for research support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Waiting for the Right Offer: Laboratory Evidence on How News Affects Bargaining
Tingting Ding and Steven F. Lehrer
NBER Working Paper No. 34543
December 2025
JEL No. C70, C92, D82

ABSTRACT

We conduct a series of laboratory experiments that implement the Daley and Green (2020) model to examine whether the gradual, exogenous revelation of sellers' private information influences the occurrence of trades in a bilateral bargaining setting with a static lemon condition. We find that while information does not increase efficiency, it reduces the likelihood that buyers incur losses when trading with low-quality sellers. Anticipating that additional signals will arrive, buyers hesitate to finalize deals immediately and exhibit a "waiting for news" effect, making sizable offer adjustments only when sufficient positive signals have accumulated. In contrast, informed sellers are less sensitive to news but appear to wait for the offer that they deem acceptable.

Tingting Ding
James Madison University
dingtx@jmu.edu

Steven F. Lehrer
Queen's University
Department of Economics
and NBER
lehrers@queensu.ca

1 Introduction

Many bargaining settings are characterized by the presence of private information that can dynamically unravel as negotiations proceed if each party receives relevant yet noisy pieces of new information. Consider, for instance, a large pharmaceutical firm's potential acquisition of a small biotech company that has new drugs under development. The biotech company initially has significant private information on the progress of each drug trial and is aware that clinical news, including the public reporting of trial results, would drastically change their company's valuation to the acquiring firm. During negotiations, the acquiring firm faces a trade-off when making an offer at any point in time, since it would be forgoing the opportunity to trade with that biotech company later when additional information arrives.

In a recent paper, Daley and Green (2020) formally explore how gradual, exogenous public revelation of sellers' private information influences the length of delay in a bilateral bargaining setting where a static lemon condition holds. The bargaining takes place between an uninformed buyer and an informed seller, who could be of high or low type. The high type seller produces high-quality products at a high cost, while the low type seller produces low-quality products at a low cost. Trade with the high type seller will never occur immediately, as the buyer's expected value of the good is lower than the cost of a high-quality product. However, prior to a deal being reached, pieces of information will continuously arrive, providing noisy signals regarding the seller's type. When a sufficiently long series of such signals are observed, the seller's type can be revealed, allowing trade to be executed with both types of sellers in equilibrium. As such, this gradually revealed information may improve bargaining efficiency. Nevertheless, this information may decrease efficiency as it creates an incentive for the buyer to wait for more information, delaying trade. Moreover, the high type seller has an incentive to enhance her bargaining position by waiting until the news clearly reveals her type, whereas the low type seller may benefit by mimicking the high type seller, resulting in a no-trade period emerging in equilibrium.

The structured bargaining model developed by Daley and Green (2020) offers tractable predictions for equilibrium behavior. Our motivation for testing this model stems from its broad relevance to real-world settings, such as mergers, acquisitions, and labor market decisions, where new information unfolds progressively during negotiations. However, empirically testing this model with field data presents significant challenges due to missing covariates; researchers often cannot accurately track how the information environment evolves throughout the bargaining process or document every offer exchanged between parties.[1] To address these challenges, we employ laboratory experiments to enhance our understanding of how gradually revealed, yet noisy, information influences bargaining dynamics. Specifically, we conduct a series of experiments with three treatments that vary both the availability and quality of information provided to subjects across sessions, in which public information progressively discloses the seller's private information and quality of information is represented by the signal-to-noise ratio.

Contrary to theoretical predictions, we do not find that incremental disclosure of information increases efficiency. In the control treatment where no information is provided, buyers increase their offers much faster than is optimal, resulting in more deals being reached. In contrast, in treatments with gradual information revelation, buyers hesitate to reach deals at the outset of bargaining and wait for additional news before tailoring offers to each seller type. Furthermore, high type sellers take longer to negotiate improved prices when they believe their type has already been revealed. These patterns become more pronounced as information quality improves. Consequently, neither the introduction of information nor its improved quality enhances efficiency.

Moreover, our findings suggest that information benefits only buyers, not either type of seller.

[1] As a result, researchers in both accounting and finance typically examine how unscheduled corporate events rather than sequences of information impact outcomes of interest. These unscheduled events include the trading of private debt (Wittenberg-Moerman, 2008), dividend payouts (Kalay, 2014), and asset write-offs (Haggard, Howe and Lynch, 2015). Further motivation for the importance of information asymmetry on market efficiency is provided by Brogaard et al. (2021), who conclude that firm-specific information revealed through public announcements accounts for 37% of the variation in stock prices over a 20-year period.

With access to information, buyers more effectively screen the seller's type: while they incur losses in 48.3% of transactions with low type sellers in the control treatment where no information is available, this proportion declines to 16.7–22.9% in treatments where information incrementally unfolds. In contrast, we observe that information reduces the payoffs of low type sellers not only because they are less able to mimic high type sellers, but also because they are more willing to accept offers that equally divide the trading surplus. While high type sellers receive more offers that exceed their cost, and these offers arrive earlier when information about their type is progressively revealed, they take longer to negotiate an improved offer, reducing their likelihood of reaching a deal. As a result, high type sellers do not gain from the introduction of information.

While our results show that both buyers and sellers strategically respond to gradually revealed information, there is a qualitative departure from the prediction of Daley and Green (2020): buyers reluctantly engage in experimentation to facilitate screening, which partially contributes to the failure of information to enhance efficiency. This experimentation involves making offers that are strictly greater than the value of low-quality goods but strictly lower than the cost of high-quality goods. If such offers are accepted, they must be accepted by low type sellers, resulting in a negative payoff for the buyer. However, experimentation can increase buyers' continuation payoffs, as a rejected offer accelerates belief updating, thereby expediting the screening process. As a result, the potential cost of transacting with a low type seller could be offset by the benefit of securing an earlier agreement with a high type seller. Under our design, we expect 96–97% of buyers' offers to align with experimentation, but only 4–6% of offers follow this pattern in the lab. This reluctance intensifies as subjects gain experience, with the rate of experimentation declining further in later games.

To understand why subjects deviate from the theoretical prediction, we consider two potential mechanisms: fairness considerations and behavioral responses to information, and assess their respective roles in shaping bargaining dynamics. We begin by incorporating fairness preferences into the model. Fairness concerns are frequently used to explain data from bargaining experiments and appear evident in our results, particularly in the treatments involving information. For instance, while buyers exhibit a "waiting-for-news" effect by making conservative offers at the beginning of bargaining, they eventually raise their offers to a level that equally divides the trading surplus for low-quality goods. As bargaining progresses, they continue adjusting their offers around this level as information gradually arrives, until they begin to believe the seller is of high type. Additionally, low type sellers are more receptive to accepting offers that equally divide the trading surplus, while high type sellers take additional rounds to negotiate an improved offer that moves toward a fairer outcome. By incorporating inequality-averse preferences into the model, we find that the predicted frequency of experimentation decreases from 96–97% to 13–17%, and the simulated bargaining dynamics align more closely with observed behavior in the lab.

While the business world is often perceived as prioritizing self-interest over fairness, bargaining patterns documented in many real-world settings involving business professionals or trade negotiators suggest otherwise (see Keniston et al., 2024). For instance, Kisgen, Qian and Song (2009) conclude that fairness opinions were prevalent in mergers and acquisitions between 1994 and 2003, with fairness opinions presented in approximately 80% of deals reached by the seller and 37% of those reached by the buyer. Similarly, Chaigneau, Edmans and Gottlieb (2025) provide a summary of survey evidence indicating that fairness considerations influence how directors and investors structure CEO pay contracts.

However, the modified model that incorporates equity considerations still cannot account for why buyers without access to additional information make more aggressive price offers and are more likely to engage in experimentation than those in treatments where information is gradually revealed. We conjecture that the behavior observed in the control treatment reflects a failure of contingency thinking, as proposed by Martínez-Marquina, Niederle and Vespa (2019), who find that buyers in a lemon market tend to make higher-than-optimal offers when facing uncertainty.

While many individuals struggle to understand the adverse selection problem even after gaining experience, the gradual reduction of uncertainty through information enables buyers to make more optimal decisions, leading to fewer losses and potentially fostering more stable markets.

To shed further light on how subjects use information and form bargaining strategies, we conduct an econometric analysis of their decisions, based on the sequence of news they receive and their matched partner's choices. The results show that the change in buyers' offers is positively correlated with the stock of news observed up to that point, providing additional empirical evidence of the "waiting-for-news" effect: buyers wait for sufficient positive information before raising their offers. Their decisions also reveal stronger responses to positive signals than to negative ones. In contrast, sellers are less responsive to news, likely because they already know their own type. Instead, their decisions appear to hinge on whether offers reach what they perceive as the "right level."

Our findings parallel many real-world bargaining contexts, such as mergers and acquisitions, where negotiations frequently stall under uncertainty and resume only after sufficient positive news arrives, a pattern consistent with the "waiting-for-news" effect (Bhagwat, Dam and Harford, 2016). Similar strategic delays arise in procurement and investment decisions, where managers often postpone action amid uncertainty surrounding tariffs, subsidies, or geopolitical risks (Barbosa et al., 2020). Finally, even within internal organizational contexts such as performance evaluations and product development reviews, managers and employees often engage in implicit bargaining over recognition, resources, or strategic direction. In these settings, decision makers tend to react more strongly to positive signals than to negative ones (e.g., Goller and Späth, 2024), a tendency that mirrors the asymmetric responses to news observed in our experiment. Taken together, these parallels strengthen our confidence that the findings may generalize beyond the laboratory, illustrating how frictions and behavioral tendencies can lead to inefficient delays, even when timely action would be optimal in the absence of such influences.

Our experiment, to the best of our knowledge, is the first to allow subjects to receive frequent, noisy information throughout the bargaining process. While previous studies (e.g., Kagel and Levin, 1986; Grosskopf, Rentschler and Sarin, 2018) explore how subjects respond to a single public signal in an incomplete information setting, real-world environments more often involve noisy information that arrives gradually rather than in a single shot. Our design captures this feature and allows us to investigate not only how subjects respond to new information, but also how they behave in anticipation of frequent information arrival. Whether such information improves trade efficiency and how it affects the allocation of the trading surplus between parties depends crucially on the strategies that subjects adopt. We suggest that data from the lab can enhance our understanding of how bargaining parties develop strategies, whether sophisticated or based on heuristics, to navigate complex adverse selection problems. These insights may subsequently inform both theoretical advancements and the design of information disclosure policies.

The rest of the paper is organized as follows. Section 2 provides a brief review of the related experimental literature. Section 3 outlines the model of Daley and Green (2020), and then describes the discrete approximation that we implement in the lab. Section 4 introduces the experimental design and implementation. Section 5 presents the results, a discussion of which is given in Section 6. A final section draws the main conclusions.

2 Related Experimental Literature

The results presented in this paper relate to experimental studies of bargaining and ultimatum games under asymmetric information, learning the optimal stopping strategy in real-options problems, and belief updating.

Roth (1995) explores early bargaining experimental studies and concludes that subjects generally incorporate both strategic considerations and fairness concerns in their decision-making. A series of bargaining experiments involving information asymmetries have demonstrated that subjects negotiate strategically. Experiments examining the monopolist's problem find that, as predicted by sequential equilibrium, the uninformed monopolist employs a monotonic sequence of prices to screen buyers whose values are independently distributed (Rapoport, Erev and Zwick, 1995; Reynolds, 2000; Cason and Reynolds, 2005).[2] Similarly, in a bilateral bargaining game where the lemon condition holds, Bochet and Siegenthaler (2018) find that uninformed buyers screen by using an increasing price sequence to exhaust the patience of informed low type sellers in treatments with discounting. Furthermore, Embrey, Frechette and Lehrer (2015) conduct bargaining experiments involving two-sided asymmetric information and find that players tend to mimic behavioral types, even when such behavior does not always align with fairness considerations. Fanning and Kloosterman (2022) assume that subjects' fairness preferences are private information and present evidence that subjects adopt strategic postures in bargaining when comparing infinite-horizon games with one-shot ultimatum games.

On the other hand, previous experimental literature shows that asymmetric information influences fairness concerns in ultimatum and bargaining games. This body of research reveals that when players have limited information about the size of the pie to be divided, proposers (buyers) tend to offer a smaller share, and recipients (sellers) are likely to accept such offers. However, as more information is revealed about the pie size, conflicts over how to split the pie increase, leading to higher rejection rates (Mitzkewitz and Nagel, 1993; Straub and Murnighan, 1995; Croson, 1996; Rapoport, Sundali and Seale, 1996; Huang, Kessler and Niederle, 2024). Similarly, Babcock, Loewenstein and Wang (1995) and Bochet, Khanna and Siegenthaler (2024) document that more information may not necessarily improve efficiency in bargaining games.

Our study aims to bridge the two branches of literature discussed above. As news gradually reduces information asymmetries, we observe a shift in the bargaining dynamics among subjects in

[2] While sellers generally follow a decreasing price path, some sellers raise their prices toward the end of the sessions, potentially reflecting retaliatory behavior.

the lab from strategic mimicking and screening to a greater emphasis on equitable surplus division. However, this heightened prominence of fairness concerns does not imply a change in subjects' preferences. Instead, the change is more likely driven by the environment, consistent with Fehr and Schmidt (1999), who note that fairness considerations are commonly observed in bilateral bargaining but tend to diminish in competitive markets. While in their setting this change is shaped by market structure, in ours it results from the reduction of asymmetric information.

Our study also relates to the literature on the optimal stopping time in real options theory, as sellers in our bilateral bargaining games need to decide when to accept the offer, an optimal termination time, as information progressively unfolds. While there is a concern that non-professional subjects may struggle to comprehend dynamic stochastic information well, previous experiments on real options problems have shown that subjects in the lab can effectively leverage continuously evolving information, and their decisions improve with the quality of information (Oprea, Friedman and Anderson, 2009; Sandria et al., 2010; Anderson, Friedman and Oprea, 2010; Descamps, Massoni and Page, 2022). In our design, subjects make their decisions based on both the evolution of exogenous information and the strategies employed by their counterparty, shedding light on how individuals leverage such information in a strategic environment.

Lastly, our analysis relates to the recent work on belief updating. A substantial body of literature in both psychology and economics suggests alternative behavioral explanations (e.g., Benjamin, 2019) for why subjects in laboratory settings seldom perfectly implement Bayes' theorem when updating their beliefs. For example, experimental subjects are often conservative, updating too little in response to signals (see Eil and Rao, 2011; Müobius et al., 2022; Coutts, 2019, among others), or neglect the correlations in information signals (Enke and Zimmermann, 2019). In our bargaining experiment, there is a theoretically established one-to-one non-linear mapping between a buyer's belief and his offer, which enables us to investigate how subjects respond to the new information signals and change their offers accordingly.

3 Summary of the Theory

3.1 Model Setup

In this subsection we provide an overview of the theoretical results for the stylized bargaining and news model of Daley and Green (2020). This is a one-sided incomplete information bargaining model in which a privately informed seller has an indivisible durable good to sell. The uninformed buyer makes frequent offers to an informed seller until his offer gets accepted by the seller.

The game is set in continuous time, denoted by t, which starts at 0 and can potentially have an infinite horizon. Both the buyer and the seller are long-lived, risk-neutral expected utility maximizers with a common discount rate $r > 0$. The seller's type $\theta \in \{L, H\}$ determines the quality of the good (low or high) for sale and is her private information. The buyer does not know the seller's type but holds a prior belief that the seller is of type H with probability $P_0 = \Pr(\theta = H)$. The value for the buyer from purchasing a type θ good is V_θ, which has a cost for the seller of K_θ. A static lemon condition holds, such that $K_H > V_L$, ensuring that there exists some non-degenerate P_0 for which $\mathbb{E}(V|P_0) < K_H$.

The model assumes that the buyer has all the bargaining power. At each time t, the buyer proposes an offer and the game ends once the seller accepts an offer. An agreed trade for price W at time t, yields payoffs to the buyer and seller respectively of $e^{-rt}(V_\theta - W)$ and $e^{-rt}(W - K_\theta)$. If an agreement is not reached, both the buyer and the seller receive a payoff of zero.

The key innovation of this model is that if a deal is not reached at time 0, both parties will start to observe a public Brownian diffusion process X that gradually reveals the seller's type.[3]

[3] If information is not revealed gradually but only provided at a single point in time, theory would make substantially different predictions. Buyers would update their beliefs immediately after receiving the singleton information, and the equilibrium collapses to Deneckere and Liang (2006), in which no information is provided, and buyers commit to a sequence of price offers to screen the seller's type. In contrast, when information progressively unfolds over time, buyers are expected to adjust price offers according to the news they receive at each point in time.

Assuming $X_0 = 0$, X_t evolves differently by seller type as follows:

$$dX_t = \mu_\theta dt + \sigma dB_t, \tag{1}$$

where $B = \{B_t, \mathcal{F}_t, 0 \leq t \leq \infty\}$ is standard Brownian motion. The informativeness of the news process is captured by the signal-to-noise ratio $\phi = (\mu_H - \mu_L)/\sigma$, where the parameters σ, μ_H, μ_L and $\mu_H > \mu_L$ are common knowledge. The news is completely uninformative if $\phi = 0$, and larger values of ϕ reveal more information about the seller's type.

To shed intuition on the equilibrium, we first consider two extreme cases without the provision of news. If θ is common knowledge and the buyer possesses all of the bargaining power, news is not needed and trades would take place immediately at $t = 0$. At the other extreme, θ is the seller's private information, but the uninformed buyer is allowed to make frequent offers over the infinite horizon. Deneckere and Liang (2006) restrict the buyer to updating his belief and preparing new offers solely based on the history of rejected offers, and show that there exists a unique sequential equilibrium. This equilibrium is characterized by a sequence of the buyer's price offers p_t. For each $t < T$, the H type seller rejects the offer, while the L type seller is indifferent between accepting the price p_t now or rejecting it and waiting for the subsequent price p_{t+1}. Further, at time T, any remaining sellers accept the offer $p_T = K_H$. Thus, trade eventually occurs for both types of sellers but following a long delay.

Daley and Green (2020) characterize the structure of the equilibrium when the buyer can learn about the seller's type θ gradually from a stochastic news process. The buyer begins the game with prior P_0, which can be updated by both the news process presented in equation (1) and whether the seller has rejected all past offers. The probability P_t that the buyer believes the seller is of H type at time t evolves according to Bayes rule:

$$P_t = \frac{P_0 f_t^H(X_t)(1 - S_{t^-}^H)}{P_0 f_t^H(X_t)(1 - S_{t^-}^H) + (1 - P_0) f_t^L(X_t)(1 - S_{t^-}^L)},$$

where S^{θ}_{t-} represents the cumulative probability that a type θ seller has accepted an offer before time t, and f^{θ}_t denotes the density of X_t conditional on θ.

In the stationary equilibrium, the buyer's belief process $P = \{P_t : 0 \leq t \leq \infty\}$ follows a time-homogeneous process and can be re-expressed using the log-odds ratio $Z_t \equiv \ln\left(\frac{P_t}{1-P_t}\right)$.[4] The equilibrium strategy can be characterized as follows: there exists a threshold β that increases with the quality of news ϕ. When $Z_t \geq \beta$, the buyer offers K_H and both types of sellers would accept it with probability 1. When $Z_t < \beta$, however, the buyer engages in a price experimentation strategy. This strategy involves making offers that are strictly greater than V_L and strictly less than K_H, which would be accepted only by an L type seller and guarantee the buyer a negative payoff.[5] The model predicts that in equilibrium, the buyer continues to engage in this strategy, as the rejection of these experimentation prices accelerates belief updating and thus the screening process.

Although the buyer has the ability to negotiate prices over time, the model predicts that he earns the same payoff as if he made no offers until belief reaches β and trade occurs at K_H. The reason lies in Coasean logic: due to his ability to learn from news, the buyer is tempted to make profitable offers to L types, behaving as if competing with his future self. If he could extract positive surplus from such trades, he could profitably deviate by offering a price just above the L type's continuation value but still below V_L. Yet such offers would, in turn, raise the continuation value of the L type. As a result, in equilibrium, the L type's continuation value must increase to the point where the buyer earns exactly the same payoff as if he had no ability to negotiate the offers. The experimentation strategy helps reduce trade delay, and all efficiency gains are captured by the L type sellers.

[4]The game has a unique stationary equilibrium, but there may exist other non-stationary equilibria. Given the nature of this complex model, the non-stationary equilibria are difficult to fully characterize.

[5]This experimentation price is given by $W(Z_t) = V_L + C(u-1)e^{uZ_t}$, where $C = \frac{K_H - V_L}{u-1}\left(\frac{u}{u-1}\frac{K_H - V_L}{V_H - K_H}\right)^{-u}$, and $u = \frac{1}{2}\left(1 + \sqrt{1 + 8r/\phi^2}\right)$. The H type seller will reject this offer, and the L type seller will accept it with intensity $1 - e^{-Q_t}$, where $Q_t = \ln\left(\frac{1-S^H_{t-}}{1-S^L_{t-}}\right)$.

When players are sufficiently patient in this model, the buyer increases (lowers) his offer upon receiving positive (negative) news. This contrasts with the equilibrium without news about the seller's type, in which buyers consistently increase their offers after each rejection.

3.2 Discrete Time Implementation

We follow the experimental literature (Oprea, Friedman and Anderson, 2009; Anderson, Friedman and Oprea, 2010), which tests continuous-time models by approximating them with rounds of extremely short duration. We use a binomial process in the lab to approximate the Brownian news process. In discrete time, news evolves in units of $t = \Delta, 2\Delta, ...$ for some $\Delta > 0$. Assume at each round t, prior to reaching an agreement, news about the seller's asset, X_t, is revealed via a binomial process. As before, we assume $X_0 = 0$, and for each seller type X_t evolves according to

- If $\theta = H$,

$$X_{t+1} = \begin{cases} X_t + \sigma\sqrt{\Delta} & \text{with prob. } q \\ X_t - \sigma\sqrt{\Delta} & \text{with prob. } 1-q \end{cases}$$

- If $\theta = L$,

$$X_{t+1} = \begin{cases} X_t + \sigma\sqrt{\Delta} & \text{with prob. } 1-q \\ X_t - \sigma\sqrt{\Delta} & \text{with prob. } q \end{cases}$$

where $q = \frac{1}{2}(1 + \frac{\mu}{\sigma}\sqrt{\Delta})$ for some $(\mu, \sigma) \in \mathbb{R}^2_{++}$ captures the quality of news ($\mu_H = \mu$, $\mu_L = -\mu$). As $\Delta \to 0$, the equilibrium converges to that of the continuous-time model.

4 Experimental Design and Testing Hypotheses

4.1 Experimental Design and Implementation

The experiments were conducted in the CESS laboratory at New York University. One hundred and forty-four students were recruited from the general student population of the university using the

CESS recruitment software. We conducted three treatments, each with four sessions, to investigate whether providing information, as well as the quality of such information, influenced subjects' behavior. Each treatment was implemented using a between-subjects design and each subject participated in only a single session. The experiment was programmed with z-tree programming software (Fischbacher, 2007). Table 1 summarizes the design parameters and number of sessions for each treatment that were carried out.

Table 1: Summary of the Design of All Laboratory Sessions

Treatment	# of Buyers	News	# of Groups per Session	# of Session	# of Subjects
C1	1	N/A	6	4	48
B1L	1	$\phi = 2$	6	4	48
B1H	1	$\phi = 4$	6	4	48
Total				12	144

Irrespective of the treatment, subjects were asked to make decisions in 12 trading games. Each subject was assigned to be either a seller or a buyer in the very first trading game, and remained in that role throughout the session. At the start of each game, buyers and sellers were randomly matched. Buyers then input an initial offer for the good. If the seller accepted this offer, a trade was made and the game ended. If the seller rejected the initial offer, the group moved to the second stage that may consist of multiple rounds.

Figures 1 and 2 present sample screenshots of the second stage in the treatments involving information. In the second stage of these treatments, a new piece of information was provided to subjects in each round prior to the deal being reached. As discussed in section 3.2, the information was presented as a tick that evolved according to a binomial process in the bottom left panel on each subject's computer screen. We set $X_0 = 0$, $\mu = 0.1$, $\Delta = 0.01$ minutes. In the treatment with low-quality information ($\phi = 2$), the tick moved up (down) with a probability of 0.55 and moved down (up) with a probability of 0.45 if the seller was of H (L) type. In the treatment with high-quality information ($\phi = 4$), the tick moved up (down) with a probability of 0.60 and moved down (up) with a probability of 0.40 if the asset was of H (L) type. To ensure that all jumps from ticks within the news process appeared similar on the subjects' screens, the axes were adjusted across treatments.

Figure 1: Screenshot for the Buyer in the Treatment Involving Information

In each second stage round, buyers had 6 seconds to make a new offer by moving the computer mouse along a slider bar on their screen. The computer treated the final position of the slider as the desired offer. If a buyer did not move the slider bar within the round, the computer assumed the buyer did not want to change his offer from the previous round. The seller also had 6 seconds to decide whether to accept the offer from the buyer, and the computer had a default decision of rejection if a seller did not reach any decision within 6 seconds.[6] If the seller accepted this offer, a

[6]Theoretically, the time interval is only 0.01 minutes, *i.e.*, 0.6 seconds. But in the lab we paused the clock for 6 second to allow subjects to react and make a decision. Previous continuous-time experiments did not need to provide a reaction time since the game under study did not involve interactions, thereby allowing the clock to keep running. The subjects in our experiments were very active. Sellers usually clicked the rejection button instead of letting the computer reject the offers for them, and 72.89% of the 9,628 buyer offers made in the second stage rounds involved

Figure 2: Screenshot for the Seller in the Treatment Involving Information

trade was made and the game ended. Otherwise, if the seller rejected the offer, the group moved to another round. This rejected offer was added to a graph in the bottom right panel of each computer screen containing the chronological path of previously rejected offers in that game. In the control treatment where no information was provided, only the path of rejected offers were provided at the base of the computer screen.

To induce the common discount rate, block random termination with a per-game probability of 0.04 was used.[7] Subjects played the above game in pre-announced blocks of 12 rounds, where only at the end of a block, were subjects told whether the game ended within that block and, if so, in what round; otherwise, they were told that the game had not ended yet and they started a new block. If a deal was reached before random termination, the type of traded asset as well as

a change from moving the slider bar.

[7] Fréchette and Yuksel (2017) present evidence that block random termination (BRT) performs similarly to random termination. Further, BRT provides an opportunity for more data collection with blocks while maintaining the same expected length of match. This feature is particularly important in our experiments that introduce noisy information, since theory predicts a long bargaining duration. BRT has been applied in numerous contexts including experimental tests of dynamic games (Vespa and Wilson, 2019) and of bond markets (Weber, Duffy and Schram, 2018).

the subject payoffs in this game were announced. The profit (or the loss) of the buyer whose offer was accepted was calculated as the difference between his value and the offer he made, while the profit (or the loss) of the seller was calculated as the difference between the amount she received and her cost. Otherwise, if the deal was not reached before the termination, both the buyer and seller earned zero.

To attempt to ensure common knowledge, the instructions provided to subjects clearly stated that in each of the 12 trading games, one third of sellers were randomly allocated to H type assets and two thirds were randomly allocated to L type assets. The value and cost of assets were fixed across all trading games and were denominated in experimental points. The H type asset was valued at 300 points for the buyer while costing the seller 180 points. The L type asset had a value of 60 points for the buyer and cost the seller 0 points. In each trading game, buyers and sellers were randomly rematched.

We selected parameters across sessions to ensure the optimal stopping time varied between 5 and 30 rounds in most treatments.[8] This would ensure that we have sufficient statistical power to detect differences across treatments when carrying out a handful of sessions in two hours, a reasonable duration for this subject pool.

In each game every subject received an endowment of 120 points to cover any possible losses. At the end of session, subjects were paid a participation fee of $10, as well as their payoff from a randomly chosen game that included any endowment.[9] The conversion rate from experimental points to U.S. dollars was 12 points per $1. The average earnings were $22.92. A typical experimental session lasted between one and a half to two hours, which included 30 minutes of instructions

[8] As documented in Table 2, the optimal stopping time in the C1 treatment is greater than 30 rounds. However, subjects in the C1 treatment reached agreements much earlier than predicted and even completed the sessions faster than those in the B1L and B1H treatments. In the results section we discuss why behavioral biases rather than time pressure led buyers to bid aggressively in the control treatment.

[9] If a subject incurred a loss exceeding 120 points in his chosen payment game, he received only his participation fee of $10. This outcome occurred for 6 subjects in the C1 treatment (approximately 4% of all subjects), but for none in the B1L or B1H sessions.

and training. The training portion allowed each subject to experience 2 practice games as a seller and 2 practice games as a buyer. In the treatments involving information, the training section also provided up to 20 examples of the Brownian news process to help subjects gain familiarity with this element of the experiment.[10]

To minimize potential confounders, instructions were previously recorded and played over the CESS audio system for subjects who interacted solely through computer terminals. While the recording was played, the text being read was projected on a screen and the subjects were provided a copy of the instructions. In each treatment involving information, we fixed a game-specific path of news for each quality type of the traded good. Thus, the path of news varied across games for each type of good, in an identical fashion across sessions. We also sought to ensure that the path of termination in each game was similar across treatments so that differences in implementation would not be confounding estimates of the treatment effects. Online Appendix A provides the instructions, as well as the paths of news for each game in the treatments involving information.

4.2 Testable Hypotheses

Our empirical investigation of the laboratory data is guided by a series of comparative static predictions from the Daley and Green (2020) model, which are summarized in Table 2.[11] While our design does not cover all possible scenarios, the parameters were deliberately chosen to detect treatment differences in predicted behavior with 90% statistical power at a 5% significance level, given the number of subjects recruited.[12] Below, we group the hypotheses by buyers, sellers, and

[10]The examples of the Brownian news process presented in each of the training games were identical across each session. The evolution of the news processes differed from those displayed in the second stage of each game that was considered for payment.

[11]The calculation of theoretical predictions for the C1 treatment is straightforward. For the B1L and B1H treatments, we computed predictions by simulating the bargaining dynamics using 6,000 randomly generated series of news, which may result in some predictions being slightly different from the closed-form solutions. For instance, there is a slight difference between the payoff prediction for the L type seller and the value of the buyer's first offer, although, in theory, these two values should be equivalent.

[12]Specifically, our hypotheses are robust to parameter choices when comparing the C1 sessions with B1L (B1H) sessions, but they may be sensitive to changes in information quality. For instance, under alternative parameter

properties of the deals reached.

Table 2: Theoretical Predictions for Each Treatment

Treatment	**C1**	**B1L**	**B1H**
First Offer	20.68	67.50	67.70
	(0.00)	(0.00)	(0.00)
Bargaining Duration			
H Type Seller	53.00	17.53	11.14
	(0.00)	(9.74)	(6.11)
L Type Seller	26.74	9.48	5.39
	(25.57)	(2.84)	(1.62)
Accepted Offer			
H Type Seller	180.00	180.00	180.00
	(0.00)	(0.00)	(0.00)
L Type Seller	97.51	100.94	84.83
	(74.10)	(22.07)	(20.86)
Payoff			
Buyer	12.44	5.01	14.23
	(21.26)	(15.93)	(14.83)
H Type Seller	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
L Type Seller	20.68	67.33	67.22
	(0.00)	(17.80)	(16.31)
Efficiency	26.23	49.89	59.04
	(17.34)	(6.79)	(6.17)
Experimentation Offers			
Fraction of offers $\in (V_L, K_H)$	0.33	0.97	0.96
	(0.23)	(0.01)	(0.01)

The model predicts that buyers engage in experimentation with information, while in the absence of information, they initially make low offers and use delay to screen the seller's type. Consequently, buyers are expected to make higher initial offers in the B1L and B1H treatments. Under our design parameters, we do not expect discernible differences in initial offers between these two treatments as the information quality improves.

Hypothesis B.1 *Buyers make higher initial offers when information is gradually revealed.*

We define an offer weakly greater than K_H as a high-type offer, since K_H is the minimum

configurations, improved information quality could potentially increase delays.

price accepted by a rational H type seller. Buyers start to propose a high-type offer when they believe that a seller is likely an H type. As both Deneckere and Liang (2006) and Daley and Green (2020) predict that buyers can utilize delay (and news) to effectively screen the seller's type, we expect that buyers are more likely to propose a high-type offer to H type sellers in each treatment. Moreover, under our design parameters, Daley and Green (2020) predict that buyers screen the seller's type faster with the aid of information, and they propose a high-type offer in earlier rounds as the information quality improves in the B1H treatment.[13]

Hypothesis B.2 *Buyers are more likely to propose a high-type offer to H type sellers than to L type sellers, and they propose a high-type offer to H type sellers in an earlier round when information is gradually revealed and its quality improves.*

Under our design parameters, 96-97% of buyers' price offers in the B1L and B1H treatments are expected to be consistent with experimentation. This is significantly higher than the corresponding fraction, 33%, in the C1 treatment.

Hypothesis B.3 *Buyers are more likely to engage in experimentation when information is gradually revealed.*

Turning to sellers, Daley and Green (2020) predict that information leads both seller types to reach deals earlier, and under our design parameters, bargaining duration decreases, and L type sellers receive lower average offers as information quality improves.

Hypothesis S.1 *L type sellers reach deals earlier when information is gradually revealed, and they accept lower offers in early rounds as the information quality improves.*

Hypothesis S.2 *H type sellers reach deals earlier when information is gradually revealed, and there is no difference in the average of the accepted offers across treatments.*

[13] In the absence of news, the buyer can use delay to extract surplus from L type sellers, gradually increasing offers up to K_H as his belief approaches the threshold $\underline{z}$ at which the buyer makes zero profit. However, with the arrival of gradual information, the buyer can no longer commit to this strategy. As learning occurs, the buyer's belief almost surely shifts away from $\underline{z}$ in a short time. This generates the Coasean force, prompting experimentation and accelerating the belief's movement towards the threshold for making a high-type offer.

Daley and Green (2020) predict that the buyer's payoff improves with the quality of information. The increased trade surplus with information, which is reflected by a shorter bargaining duration, will be allocated to L type sellers. H type sellers do not benefit from information, since they have no bargaining power and are expected to receive a payoff of zero in all three treatments.

Hypothesis P.1 *L type sellers benefit from information, whereas H type sellers do not.*

Hypothesis P.2 *In the sessions involving information, buyer payoffs increase as the quality of information improves.*

Buyers' experimentation accelerates screening in the information treatments, thereby reducing inefficiency due to delay. Efficiency is highest in the B1H treatment and lowest in the C1 treatment.

Hypothesis E *Efficiency increases when information is gradually revealed and the information quality improves.*

5 Results

Table 3 summarizes the outcomes observed in the lab for each treatment. The analysis in this section will focus on data from the last 6 trading games.[14] The p-values reported throughout this section are calculated using random-effects regressions with a complete set of treatment indicators, with standard errors clustered at the session level, except for efficiency, which is calculated for each game. Additionally, we use Heckman's procedure to estimate parameters in linear models when dealing with censored data.

[14]We focus on the last 6 games as the earlier trading games may involve subjects familiarizing themselves with the environment and exhibit greater volatility in decision-making. This practice is common in the literature of dynamic games (Dal Bó and Fréchette, 2018). Additionally, in the C1 treatment, we exclude two specific subjects and their group partner from the analysis, as these subjects displayed significantly different behavioral patterns compared to other participants. In online Appendix C, we demonstrate robustness to including these subjects and to using all games, finding no qualitative differences from the results reported in Table 3.

Table 3: Descriptive Statistics in the Last 6 Trading Games

Treatment		C1	B1L	B1H
First Offer		27.47	25.13	18.27
		(21.99)	(21.63)	(10.83)
Accepted Offer				
	H Type Seller	187.71	213.74	208.08
		(58.15)	(14.49)	(25.92)
	L Type Seller	109.13	71.63	60.87
		(67.99)	(63.38)	(63.96)
Payoffs				
	Buyer	−9.67	5.08	15.46
		(79.89)	(56.79)	(60.79)
	H Type Seller	3.77	13.35	14.63
		(40.32)	(18.93)	(23.33)
	L Type Seller	90.73	58.20	52.63
		(74.31)	(63.60)	(63.00)
Deal Round				
	H Type Seller	13.38	17.16	18.92
		(9.30)	(10.74)	(10.77)
	L Type Seller	9.76	8.38	7.19
		(7.74)	(7.39)	(5.79)
Fraction of Deals Reached				
	H Type Seller	0.49	0.40	0.52
		(0.51)	(0.49)	(0.50)
	L Type Seller	0.83	0.81	0.86
		(0.37)	(0.39)	(0.34)
Efficiency		36.47	32.46	36.84
		(16.01)	(11.00)	(11.23)
Experimentation Offers		0.14	0.06	0.04
Fraction of offers $\in (V_L, K_H)$		(0.34)	(0.23)	(0.19)

Note: Each entry presents the mean and standard deviation of the row variable for the treatment denoted by the column heading. The variables deal round and accepted offer are recorded for the sample in which deals were reached.

5.1 Buyers' Behavior

Result B.1 Buyers make lower initial offers when they anticipate that new information will gradually arrive, with the most conservative initial offers observed in the B1H treatment.

Our data reject Hypothesis B.1. We find that, in contrast to the theoretical prediction, buyers' initial offers decrease from 27.47 to 25.13 and 18.27 across the C1, B1L, and B1H treatments. While the information quality is expected to have little impact on the magnitude of initial offers given our design parameters, we find that buyers make significantly lower offers in the B1H treatment ($p < 0.005$ between the B1H and C1/B1L treatments).

Figure 3: Initial Offers Proposed by Buyers

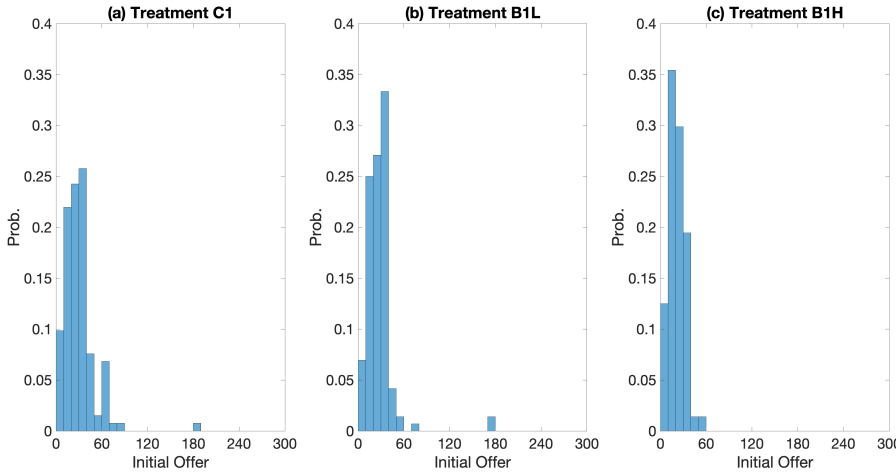

Figure 3 presents the relative frequency of initial offers across treatments. Rather than engaging with experimentation, the majority of buyers make initial offers below $V_L/2$. More specifically, 31.25% of initial offers in the B1H treatment are below 10, whereas the corresponding fractions are 24.24% and 24.31% in the C1 and B1L treatments. Given that very few sellers would accept an offer below 10, we interpret these conservative initial offers as a sign of waiting-to-learn.[15] Buyers appear hesitant to seal the deal in the first stage, and waiting-to-learn appears most pronounced in the high-quality information treatment.

Result B.2 Buyers are more likely to propose a high-type offer to H type sellers than to L type sellers in all three treatments. However, the mere introduction of information does not accelerate buyers' screening. Only when the quality of information improves do buyers make high-type offers in earlier rounds.

Our data provide partial support for Hypothesis B.2. Figure 4(a) presents the cumulative distribution function of the number of rounds until the buyer first proposes a high-type offer.[16] In

[15]Only once did two sellers in the C1 treatment accept an offer of 10, and all the remaining sellers rejected offers weakly less than 10.

[16]Notice that none of these cumulative probabilities reaches 1, since some buyers never make a high-type offer before either a deal is reached or the game is randomly terminated.

Figure 4: The Round When the High-Type Offer is Made and When a Deal is Reached

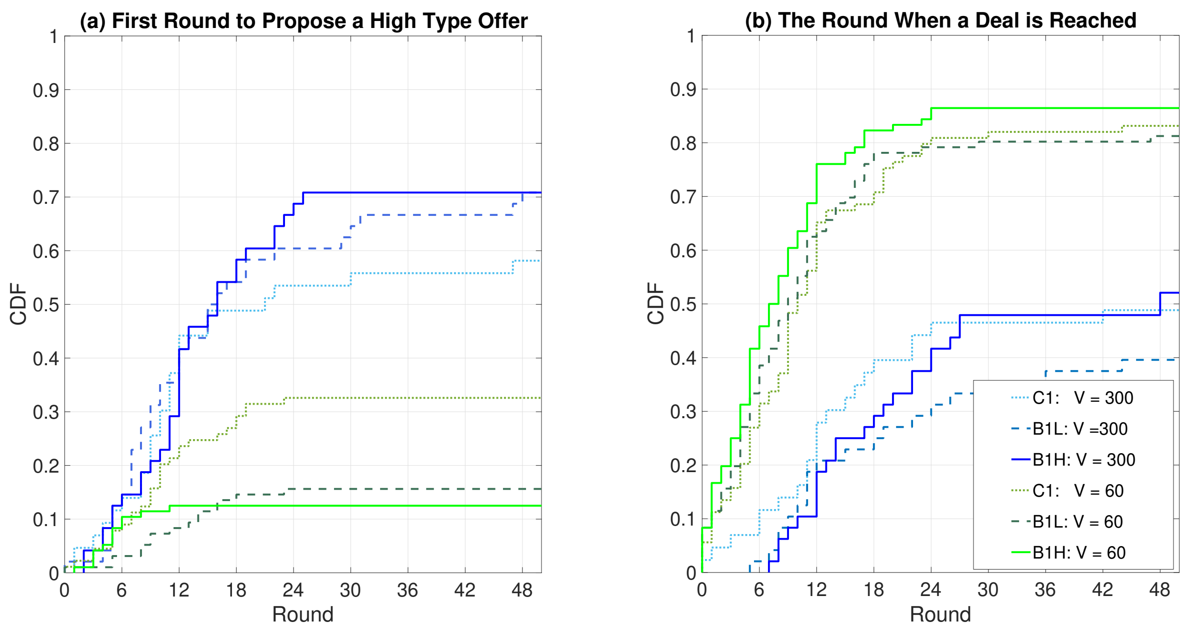

the C1 sessions, 58.1% of H type sellers and 32.6% of L type sellers received at least one high-type offer, compared to 70.8% and 15.6% in the B1L sessions, and 70.9% and 12.5% in the B1H sessions, respectively. Buyers appear to screen seller types more effectively with the aid of gradually arriving information, as they are less likely to make high-type offers to L type sellers.[17] In the B1H sessions, no high-type offers were made to L type sellers after round 10, even though 36.5% of potential L type deals remained open at that time.

We further compare the timing of buyers making their first high-type offer across treatments by estimating regression models that account for selective continuation using a Heckman correction. The results reveal that buyers make a high-type offer 3.28 rounds earlier in the B1H relative to the B1L treatment ($p = 0.05$), but we do not find a statistically significant difference between the C1 and B1L (B1H) treatments. This suggests that the mere introduction of information does not

[17]There is no significant difference in the fraction of deals reached by L type sellers across treatments. Thus, the lower fraction of high-type offers made to L type sellers in the B1L and B1H sessions is unlikely to result from their previously accepting an offer below K_H.

motivate buyers to put forth a high-type offer earlier, but higher-quality information appears to improve the pace of screening.

Result B.3 Buyers are less likely to engage in experimentation when information is gradually revealed.

Our data reject Hypothesis B.3. In contrast to theory, we observe that buyers are less likely to engage in experimentation when information is provided. The fractions of offers consistent with experimentation are 14%, 6%, and 4%, in the respective C1, B1L, and B1H treatments. Information leads to a statistically significant decline with $p < 0.005$ between the C1 and B1L/B1H treatments.

5.2 Sellers' Behavior

Result S.1 *L type sellers in sessions with information tend to accept lower offers and do so in earlier rounds. As information gradually reveals their type, they are more likely to accept offers that equally split the trading surplus, with this pattern most pronounced in the B1H sessions.*

Our data only provide evidence for shorter bargaining duration in the B1H treatment. Regressions that account for selective continuation using a Heckman correction reveal that L type sellers reach deals 2.56 rounds earlier ($p < 0.05$) and accept offers 48.25 points lower ($p < 0.001$) in the B1H treatment compared to the C1 treatment. However, these regressions do not find a significant difference in the number of rounds taken for an L type seller to reach a deal between the C1 and B1L treatment, and only reveal that buyers accept an offer of 37.49 points lower ($p < 0.001$) in the B1L treatment.

It is noteworthy that in the treatments involving information, L type sellers accept significantly lower offers than theoretical predictions ($p < 0.001$ for both B1L and B1H treatments). Upon examining the bottom row of Figure 5, which displays the distribution of offers accepted by L type sellers in each treatment, we observe that 30.12% of accepted offers in the B1H treatment

Figure 5: The Distribution of Offers Accepted by Sellers Across Treatments

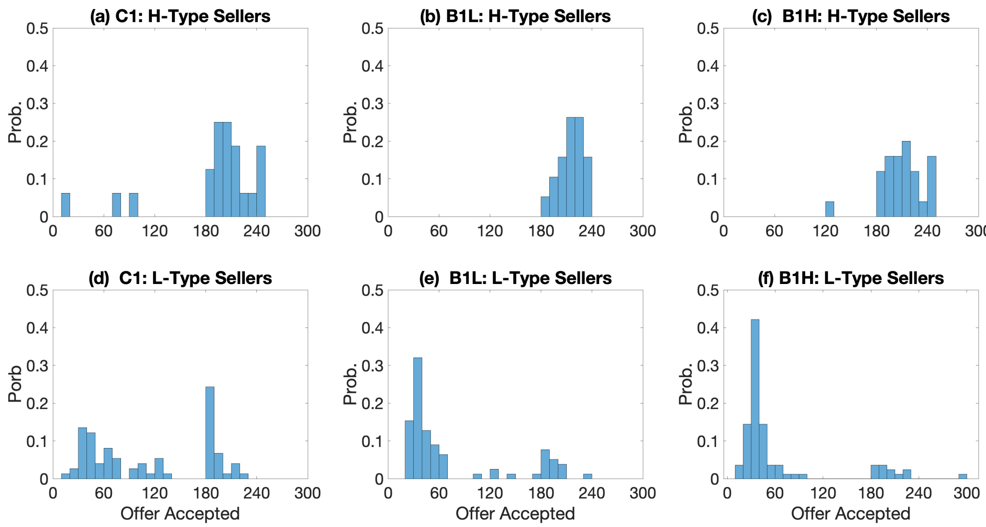

Note: The horizontal axis in each figure corresponds to the amount of the accepted offer if a deal is reached.

are exactly 30 ($V_L/2$), and 48.19% of accepted offers lie between 25 and 35. Similarly, in the B1L treatment, 37.18% of accepted offers lie within this range. However, in the C1 treatment, only 8.11% of accepted offers fall within the 25 to 35 range. As information gradually reveals their type, L type sellers tend to accept offers that equally divide the surplus and do so in earlier rounds. This behavioral pattern becomes more pronounced as the quality of information improves.

The distribution of accepted offers by L type sellers in the C1 treatment also reflects buyers' ineffective screening. In the C1 treatment, 31.46% of L type sellers successfully mimic the H type seller and reach a deal at a price greater than K_H, while the corresponding fractions in the B1L and B1H treatments are only 14.58% and 12.50%, respectively. This finding confirms the previous finding in Hypothesis B.2; buyers screen sellers less effectively in the absence of information.

Result S.2 *H type sellers take additional rounds to negotiate improved price offers in the treatments with information. As a result, there is no significant difference in overall bargaining durations across treatments, even though H type sellers receive high-type offers in earlier rounds in the B1L and B1H treatments.*

Our data reject Hypothesis S.2. Despite a 12% increase in the fraction of buyers offering a high-type price before random termination in the B1L and B1H treatments, H type sellers do not exhibit a greater likelihood of securing a deal in these treatments. This pattern runs counter to theoretical predictions. As shown in Figure 4(b), H type sellers generally bargain for a longer duration to achieve an improved offer when information gradually signals their type. Supporting this, regressions that adjust for random termination with a Heckman correction reveal that H type sellers take an additional 3.77 rounds in the B1L treatment ($p < 0.05$) and 5.53 rounds in the B1H treatment ($p < 0.10$) to reach a deal relative to the C1 treatment. Further evidence from the top row of Figure 5 shows the distribution of accepted offers by H type sellers across treatments, illustrating how information helps them achieve higher offers that move closer to an even division of the surplus. However, this comes at the cost of a lower likelihood of reaching agreements.

5.3 Trade Surplus and Efficiency

Result P.1 Information does not significantly impact the payoffs of H type sellers but reduces the payoffs of L type sellers.

Our data reject the hypothesis for L type sellers but support it for H type sellers. Contrary to theoretical predictions, L type sellers do not gain an advantage from information. Their payoffs decrease from 90.73 in C1 to 58.20 in B1L and 52.63 in B1H. This is partly because buyers do not engage in experimentation as frequently as the theory predicts, and partly because L type sellers tend to accept an offer that equally divides the surplus as information gradually arrives.

Moreover, information does not significantly improve the payoffs of H type sellers. While these sellers manage to achieve an improved offer, they do so at the expense of increased bargaining duration, resulting in discounted payoffs.

Result P.2 Buyer payoffs increase when information is gradually revealed, but information quality does not lead to a significant difference in these payoffs.

Our data does not provide statistical support for Hypothesis P.2. On average, buyers' payoffs increase from -9.67 to 5.08 and 15.46 across the C1, B1L, and B1H treatments. The differences between C1 and both B1L and B1H are statistically significant ($p < 0.05$ and $p < 0.01$, respectively). This increase is driven by both better screening by buyers in the B1L and B1H treatments, as they more accurately identify L type sellers, and by the increased willingness of L type sellers to accept offers that equally divide the surplus. However, the difference in buyer payoffs between the B1L and B1H sessions is not statistically significant ($p = 0.106$). Although the mean buyer payoff is higher in B1H, the greater variance in payoffs reduces the statistical significance of the estimated difference.

Result E Neither introducing information nor increasing the quality of information improves efficiency.

To compute the efficiency observed in the laboratory, we pool all trades with the H and L type sellers in each game and calculate the weighted average discounted surplus using a random effects regression by treatment.

Our data reject Hypothesis E. On average, we do not observe that information improves efficiency. The efficiency for the C1, B1L, and B1H treatments is 36.47, 32.46, and 36.84, respectively. The limited impact of information is due to two reasons. First, the efficiency achieved in the B1L and B1H treatments is lower than predicted by the theory. Buyers in these two treatments display signs of a waiting-for-news effect and hesitate to engage in experimentation. Moreover, H type sellers often reject offers that are only slightly above K_H and take additional rounds to demand an improved price, thereby resulting in an efficiency loss. Second, buyers in the C1 treatment appear eager to reach a deal, as they increase their offers at a faster pace than is optimal, resulting in more deals reached than the theory predicts.

6 Discussion

Our examination of Daley and Green (2020) yields mixed results. Consistent with theoretical predictions, we find that both buyers and sellers strategically respond to the gradually revealed information. Buyers screen seller types more effectively using news. Meanwhile, L type sellers attempt to mimic H type sellers when information is absent and accept offers in earlier rounds as information progressively arrives.

Nonetheless, we also observe several deviations from the model's predictions. In the treatment involving information, buyers exhibit a waiting-for-news effect, making extremely conservative offers at the start of bargaining. Additionally, fairness concerns become apparent as information gradually reveals the seller's type. These concerns facilitate agreements with L type sellers, who are more inclined to accept the social norm of evenly splitting the trading surplus. In contrast, fairness considerations hinder agreements with H type sellers, who take longer to negotiate an improved offer, since they expect a larger share once they believe buyers have identified their type. Finally, buyers reluctantly engage in experimentation, leading to lower trade efficiency but higher buyer payoffs compared to the model predictions.

In the environment modeled by Deneckere and Liang (2006), where information is absent, we find that buyers struggle to commit to waiting to screen the seller's type. Instead, their aggressive bidding behavior leads to offers that escalate too quickly, resulting in more deals than predicted and negative average payoffs due to significant losses in transactions with L type sellers. As a result, several hypotheses about the role of information are undermined.

In this section, we discuss how our experimental findings can inform directions to modify the theory to better explain the observed data. We begin by extending Daley and Green (2020) to incorporate fairness concerns into subjects' utility functions. Using simulations of the resulting structural model, we find that deviations from the experimental results diminish markedly. We next examine how buyers change their offers between successive rounds, and compare their screening behavior with the model predictions. Last, we conduct an econometric analysis focusing on the trajectory of offers and examine how buyers and sellers in the laboratory make their decisions based on the sequence of news they receive and their matched partner's choices.

6.1 Incorporating Fairness Concerns

Fairness concerns are frequently evoked when interpreting data from bargaining experiments and appear evident in our study, particularly as information is gradually revealed. Beyond the fairness considerations exhibited by sellers, we find that buyers typically set their offers to approximately split the trading surplus of low-quality goods equally, following the more conservative "waiting-for-news" offers made in the first few rounds. They then continue adjusting their offers within the range of $V_L/2$ and V_L until they begin to believe the seller is of high type. Therefore, we conjecture that the low frequency of the experimentation strategy may be partially driven by fairness considerations.[18]

We incorporate one-sided inequality aversion into the model, as subjects in our lab display a concern about earning less, but not more than, their matched partner. The payoffs are modified such that if a deal is reached at a price offer W in round t, the payoff of θ type seller is $e^{-rt}\big(x_\theta - \gamma \max\{x_B - x_\theta, 0\}\big)$, and the buyer's payoff is $e^{-rt}\mathbb{E}_\theta\big[x_B - \gamma \max\{x_\theta - x_B, 0\}\big]$, where $x_\theta = W - K_\theta$ and $x_B = V_\theta - W$. The parameter γ measures the subject's sensitivity to inequality aversion. If $\gamma = 0$, the model degenerates to Daley and Green (2020). Incorporating news to the model with

[18] An alternative explanation for buyers' low frequency of experimentation could be loss aversion. However, it does not account for why buyers' offers fluctuate around $V_L/2$ until they believe the seller is of the H type. Moreover, it fails to explain sellers' behavior, as our experimental design ensures that sellers know their costs and are unlikely to incur losses. The fairness considerations introduced in this subsection provide a coherent framework that explains both buyer and seller behavior and serves as a foundation for structural estimation and simulation.

these payoffs yields an equilibrium that shares the same structure as Daley and Green (2020).

The equilibrium is formally derived in the Appendix and is characterized by a belief threshold β^γ such that: when the buyer's belief is greater than this threshold, the buyer makes the offer $\overline{W}^\gamma$, which both types of sellers immediately accept; otherwise, the buyer offers a price strictly less than $\overline{W}^\gamma$, and only L type sellers would accept it with positive probability. Several important changes in the bargaining dynamics are worth highlighting. First, since $\overline{W}^\gamma$ is greater than K_H, the buyer will receive more negative utility if he misidentifies an L type seller. In such cases, the buyer incurs both a loss $V_L - \overline{W}^\gamma$ and an inequality cost $\gamma\big(2\overline{W}^\gamma - V_L\big)$. Consequently, buyers become more cautious in screening sellers, resulting in a longer time period before offering $\overline{W}^\gamma$. Second, inequality aversion leads buyers to make a lower initial offer, since in equilibrium L type sellers are indifferent between accepting the initial offer and waiting additional rounds until $\overline{W}^\gamma$ is offered. As such, if the initial offer is lower than V_L, we expect fewer offers consistent with experimentation.

To estimate the subject's sensitivity to inequality aversion, we assume that there is no parameter heterogeneity across subjects. In the equilibrium, the minimal price offer that an H type seller would accept is $\overline{W}^\gamma = \frac{K_H+\gamma(V_H+K_H)}{1+2\gamma}$. However, in the lab, we only observe W_i^* if a deal is reached, and random termination leads to censored data. Thus, we assume that in the lab an H type seller i accepts an offer given as

$$W_i^* = \overline{W}^\gamma + \epsilon_i$$

where $\epsilon_i \sim N(0, \sigma_\epsilon)$. This assumption allows us to express the log-likelihood function as

$$\ln L = \sum_{D_i=1} -\frac{1}{2}\Big[\ln(2\pi) + \ln\sigma_\epsilon^2 + \frac{\big(W_{i\tau} - \overline{W}^\gamma\big)^2}{\sigma_\epsilon^2}\Big] + \sum_{D_i=0} \ln\Big[1 - \Phi\Big(\frac{\max_t\{W_{it}\} - \overline{W}^\gamma}{\sigma_\epsilon}\Big)\Big]$$

where $D_i = 1$ indicates that the deal is reached in round τ at price $W_{i\tau}$, while $D_i = 0$ indicates that the deal is not reached since all past offers were below W_i^*.

We estimate the inequality aversion in each information treatment, since the results in the previous section imply that subjects may have a different perception on fairness with information or changes in information quality. Estimates of γ for each treatment are presented in the top panel of Table 4.

Using the estimated subject's sensitivity to inequality aversion, we next simulate bargaining dynamics for each treatment. The results are presented in the bottom panel of Table 4. When incorporating inequality aversion, the predicted fraction of buyers adopting the experimentation strategy decreases substantially from 97% (98%) to 13% (17%) in the B1H (B1L) treatment, which more closely aligns with the outcomes observed in the lab.[19] Furthermore, simulations of the modified model more accurately capture the bargaining durations of H type sellers and the fraction of deals they reach.

Nevertheless, we note that the modified model remains unable to explain why buyers make conservative initial offers in the sessions involving information, which also partially contributes to the extended bargaining duration for L type sellers, whereas they are more inclined to engage in experimentation when information is absent. We conjecture that the aggressive bidding observed in the C1 treatment cannot be attributed to fairness considerations. Instead, it is more plausibly driven by a failure of contingency thinking. Similarly, the waiting-for-news effect is not a result of equity considerations, but rather may reflect a form of regret aversion, since subjects try to avoid prediction errors as information is gradually revealed. Finally, this modified model cannot explain some of the abrupt offer jumps made by buyers, as its stationary equilibrium, like that of Daley and Green (2020), predicts a smooth and continuous offer trajectory. In the next two subsections, we further investigate how subjects respond to information and whether their behavioral patterns can help reconcile the gap between our findings and theoretical predictions.

[19] This modified model also predicts that buyers engage in an experimentation strategy before ultimately offering $\bar{W}^{\gamma}$. These experimentation offers are accepted only by L type sellers and guarantee negative utility to the buyer. Specifically, based on the estimated sensitivity to inequality aversion, the range of offers consistent with this experimentation spans from 30.45 to 219.57 in the B1L sessions, with 31.42% of buyer offers falling within it. In the B1H sessions, the corresponding range is from 35.93 to 214.31, with 11.68% of offers falling within this range.

Table 4: Estimates and Simulation of the Model Incorporated with Inequality Aversion

Treatment	**C1**	**B1L**	**B1H**
Estimate of Inequality Aversion Coefficient			
γ	0.50	0.97	0.67
	(0.21)	(0.36)	(0.25)
Comparing the Results from Simulations of the Model with Laboratory Data			
First Offer	**20.68**	**67.50**	**67.70**
Lab	27.47	25.13	18.27
	(21.99)	(21.63)	(10.83)
Model incorporated with γ	11.57	40.48	44.19
	(1.06)	(2.15)	(2.66)
Offer Accepted by H **Type Seller**	**180.00**	**180.00**	**180.00**
Lab	187.71	213.74	208.08
	(58.15)	(14.49)	(25.92)
Model incorporated with γ	209.85	219.57	214.31
	(5.03)	(4.01)	(4.46)
Offer Accepted by L **Type Seller**	**21.48**	**83.55**	**83.66**
Lab	109.13	71.63	60.87
	(67.99)	(63.38)	(63.96)
Model incorporated with γ	11.61	45.17	56.48
	(1.07)	(3.21)	(4.19)
Deal Round with H **Type Seller**	**26.00**	**14.00**	**9.96**
Lab	13.38	17.16	18.92
	(9.30)	(10.74)	(10.77)
Model incorporated with γ	26.00	19.12	13.38
	(0.00)	(0.17)	(0.49)
Deal Round with L **Type Seller**	**0.49**	**6.74**	**4.38**
Lab	9.76	8.38	7.19
	(7.74)	(7.39)	(5.79)
Model incorporated with γ	0.09	2.12	2.34
	(0.07)	(0.48)	(0.34)
Fraction of Deals Reached with H **Type Seller**	**0.00**	**0.75**	**0.71**
Lab	0.49	0.40	0.52
	(0.51)	(0.49)	(0.50)
Model incorporated with γ	0.00	0.42	0.67
	(0.00)	(0.01)	(0.04)
Fraction of Experimentation Offers	**0.00**	**0.98**	**0.97**
Lab	0.14	0.06	0.04
	(0.34)	(0.23)	(0.19)
Model incorporated with γ	0.00	0.17	0.13
	(0.00)	(0.04)	(0.03)

Note: Bootstrap standard deviations in parentheses. The numbers presented in bold indicate the ex-post theoretical prediction of Daley and Green (2020) given the generated news and random termination in the lab. For the policy simulations, the bootstrap standard deviations are calculated using estimates obtained from the simulated model with 6,000 bootstrap samples.

6.2 Examining How Buyers Change Their Offers Between Successive Rounds

In this subsection, we examine how buyers change their offers between successive rounds. Figure 6 provides a visual representation of these adjustments, and the color of the circle reflects the type of good for sale. Notable differences emerge across treatments: in the C1 treatment, buyers' offers are more dispersed, with many falling within a range consistent with experimentation. In contrast, in the B1L and B1H treatments, fewer offers fall within that range, and buyers appear more likely to increase their offers directly from below V_L to a level equal to or above K_H.

Figure 6: How Buyers Change Their Offers Between Successive Rounds

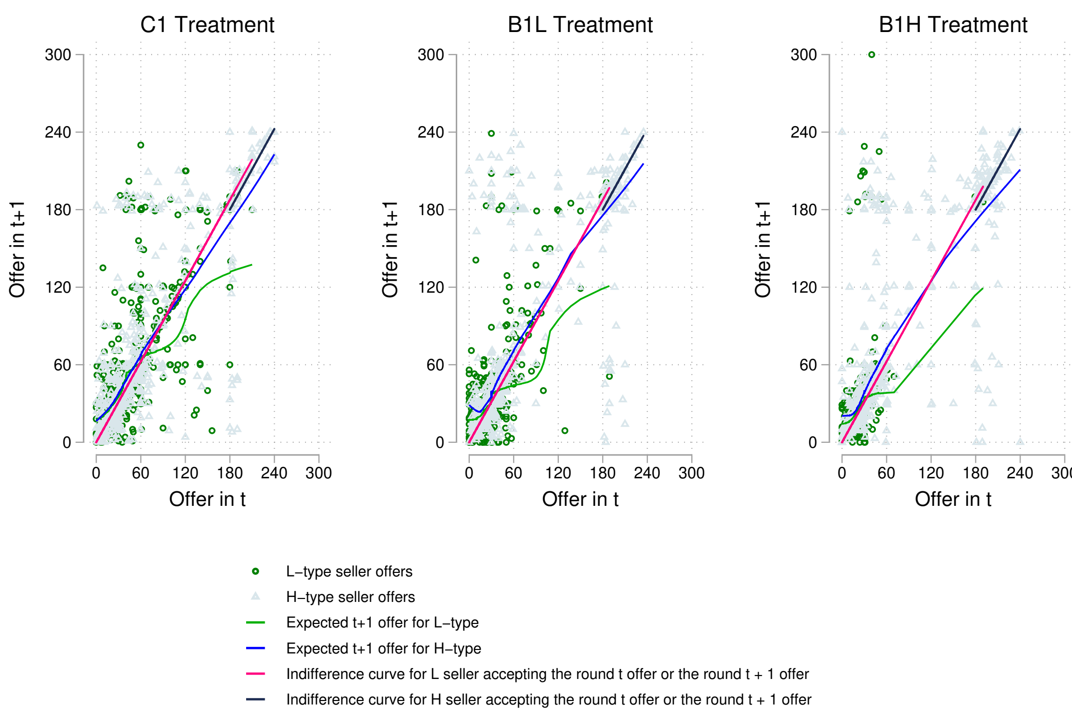

Furthermore, we include a pink curve representing the offer level at which a seller is indifferent between accepting in round t and waiting for round $t+1$. We also plot the empirical expected round $t+1$ offers as a function of the round t offer, shown separately for L type sellers (green) and

H type sellers (blue).[20] For L type sellers, their empirical curves lie above the indifference curve at lower offer levels but fall below it beyond a certain threshold. The crossing points occur at 72, 31, and 24 in the C1, B1L, and B1H sessions, respectively. When offers are below the crossing point, L type sellers are better off waiting for higher offers, while when offers are above the crossing point, they are better off accepting them immediately. Since the crossing points in the B1L and B1H sessions are close to the amount that equally divides the trading surplus, this helps explain why L type sellers are more likely to accept these near–equal-division offers in those sessions.

We also plot the indifference curve for H type sellers, shown in navy blue and defined only for the range above K_H. This curve consistently lies above the empirical offers, suggesting that once a high-type offer is made, H type sellers are better off accepting it immediately. When such offers are rejected, buyers often hesitate to increase their offers and may even lower them below V_L in retaliation. This behavior causes H type sellers to struggle to negotiate favorable prices and sometimes jeopardizes their chance of securing a deal.

We then investigate how buyers adjust their offers in the B1L and B1H treatments in response to tick movements. Under our design parameters, theory predicts that buyers should adjust their offers in the same direction as the news tick. However, we find that 27% of the time in the B1H sessions and 38% in the B1L sessions, buyers continue to offer the same price after receiving an up tick. The corresponding fractions are 28% and 34% following a down tick. This behavioral inertia aligns with the literature summarized by Benjamin (2019), which finds that subjects tend to underweight new signals, particularly if they anticipate receiving additional information. That said, buyers in the high-quality information sessions appear more responsive to new information.

Additionally, in both the B1L and B1H treatments, we observe an asymmetric response to up and down ticks: while buyers increase their offers 57% of the time in the B1H sessions and

[20] An alternative method to compute the empirical expected offer in round $t+1$ is to take the weighted average of round $t+1$ prices, with weights given by each seller type's probability of an up or down tick. This approach assumes that the round t offer fully summarizes the buyer's belief, but as shown in Section 6.3, buyers also rely on the accumulated stock of news when adjusting offers.

48% in the B1L sessions following an up tick, they lower their offers only 23% and 25% of the time, respectively, in response to a down tick. We conjecture that this asymmetric response, along with behavioral inertia, contributes to the remaining discrepancies between observed laboratory behavior and theoretical predictions. However, these findings are based on a simple unconditional relationship between the changes in offers and the direction of the most recent tick. In the next subsection, we present an econometric analysis that incorporates the matched partner's decisions and accounts for the random termination implemented in the lab.

6.3 An Empirical Understanding of Offer Dynamics and Sellers' Acceptance Decisions

To better understand how exogenous news affects offer trajectories, buyer learning, and seller acceptance decisions, we conduct econometric analysis that exploits the random variation in the information sequences observed by buyers and sellers across rounds and games, as shown in Appendix Figures A.3 and A.4. The main empirical challenge arises from the choice-based nature of the laboratory data, since buyer i in round r of game g can make an offer W_{irg} after the initial round only if the seller rejected the prior offer and the game was not randomly terminated.

To overcome this empirical challenge, we assume that the decision to continue is due to exogenous and endogenous observables that are available to the seller prior to the decision in each round. We define $D_{i(r+1)g} = 1$ to indicate that seller i reaches a deal in round r, and $D_{i(r+1)g} = 0$ if the deal is not reached. If the selective continuation is based on observables including block termination, unbiased and consistent estimates of the estimating equation for the buyer's offer can be obtained by inverse probability reweighting.[21] To estimate the probability of remaining in the

[21]That is, the population density $f(W_{irg}|I_i)$ conditional on individual characteristics I_i can be inferred from $g(W_{irg}|I_i, D_{irg} = 0)$, even though W_{irg} is observed only if $D_{irg} = 0$. Since reaching a deal in the game is an absorbing state, the weights used in each subsequent round r are simply the product of all current and past estimated probabilities, where all of the estimated probabilities for not reaching a deal in round r are obtained from a logit regression using all subjects for round $r - 1$.

sample, we assume that the seller's decision to accept the offer on the floor is a function of the information that is available to them, including the current offer, linear round effects, information from the news process, and demographic characteristics. Furthermore, since the seller is informed of the good for sale, we allow for differential responses to these covariates based on the seller's type.

Table 5: Factors Influencing Whether a Seller Accepts a Deal in Round t

Treatment	**C1**	**B1L**	**B1H**
Up Tick		-0.28	-0.12
		(0.24)	(0.29)
News Stock		-4.43	-6.44
		(3.49)	(17.50)
Round	-0.01	0.01	-0.00
	(0.02)	(0.01)	(0.02)
Offer	0.02^{***}	0.03^{***}	0.04^{***}
	(0.002)	(0.004)	(0.02)
H Type Good	-1.61	-11.49^{***}	-8.69^{***}
	(1.02)	(2.78)	(2.91)
H Type $\times$ Up Tick		1.93^{***}	0.20
		(0.53)	(0.41)
H Type $\times$ News Stock		-1.13	-30.06
		(4.35)	(33.70)
H Type $\times$ Round	-0.04^{**}	-0.05^{***}	0.08^{***}
	(0.02)	(0.01)	(0.01)
H Type $\times$ Offer	0.00	0.04^{**}	0.01
	(0.004)	(0.01)	(0.02)
Constant	-4.26^{***}	-3.62^{***}	-4.15^{***}
	(0.60)	(0.29)	(0.69)
Control	Yes	Yes	Yes
Observations	1,889	1,844	1,706

Note: Standard errors clustered at the session level in parentheses. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$.

Each column of Table 5 presents the estimates from a logit model of seller acceptance decisions. Interestingly, we find that sellers respond significantly to offers, but typically not to the news. The sellers' insensitivity to news is somewhat surprising, given the behavioral differences across the sessions with gradually arriving information and those without news. We conjecture that, since sellers are already informed of their own type, they are less responsive to gradually revealed news and more influenced by buyers' strategies, which change markedly when information is introduced. As shown in Figure 6, buyers in both the B1L and B1H sessions appear reluctant to raise their offers above $V_L/2$ unless the news reveals that the seller is of the H type. Anticipating this, L type sellers have little incentive to wait and therefore tend to accept such "fair" offers in early rounds.

Regarding H type sellers, Table 5 shows that their likelihood of reaching a deal decreases as they anticipate higher prices to offset their costs. Notably, only in B1L are the interactions between H type status and offers, and between H type status and receiving an uptick, statistically significant. This suggests a nonlinear response in seller decision-making to offer changes in those sessions. Generally, informed sellers are unaffected by news, except when they can potentially exploit buyer reactions to upticks in the noisier B1L setting.[22]

We next consider a reduced form specification to study how W_{irg} changes between rounds:

$$\Delta W_{irg} = \alpha_1 Up_{irg} + \alpha_2 News_{irg} + \alpha_3 W_{i(r-1)g} + \alpha_4 r_r + \alpha_5 I_i + \varepsilon_{irg} \tag{2}$$

where Up_{irg} and $News_{irg}$ are indicators for the most recent up tick and the level of the news stock, respectively; $W_{i(r-1)g}$ is the prior offer made in round $r-1$; r_r captures the round effect; I_i is a vector of individual characteristics that include gender and major; and ε_{irg} is a mean-zero independently distributed random shock.

Table 6 presents estimates from equation (2) that are reweighted for selective continuation, with each column corresponding to a different treatment. We find that buyers in the B1H sessions significantly increase their offers in response to both a higher cumulative stock of news and the most recent up tick. In contrast, buyers in the B1L sessions respond significantly only to the stock of news and are more likely to increase their offers as the number of rounds increases. Under the stationary equilibrium assumption, the price offer W_t is a sufficient statistic for the news history up to round t. Therefore, the stock of news should have no additional predictive power on the change in prices. The positive and significant coefficient on the stock of news suggests the waiting-for-news effect: instead of adjusting their offers smoothly, buyers wait until enough positive news has accumulated and then make a sizable and sharp jump to a high-type offer. Across all three

[22]We also find a small but statistically significant coefficient on the interaction between H type and rounds when clustering at the session level. However, this significance appears to be driven by a negative intraclass correlation, and the result is statistically insignificant when using either uncorrected or heteroskedasticity-robust standard errors.

Table 6: Factors Influencing How Buyers Adjust Their Offers Across Rounds

Treatment	**C1**	**B1L**	**B1H**
Up Tick		1.39	3.99^{**}
		(2.75)	(1.09)
News Stock		171.57^{**}	208.93^{**}
		(41.32)	(52.25)
Prior Offer	-0.12^{*}	-0.17^{**}	-0.14^{**}
	(0.05)	(0.04)	(0.03)
Round	0.17	0.45^{**}	0.23
	(0.08)	(0.09)	(0.25)
Constant	13.36^{***}	7.24^{*}	7.83^{**}
	(0.71)	(2.81)	(1.63)
Control	Yes	Yes	Yes
Observations	1,745	1,844	1,706
R-squared	0.03	0.07	0.04

Note: Robust standard errors in parentheses clustered at the session level. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$. The specifications include the full set of subject demographic characteristics, assume a homogeneous effect of each parameter in the offer equation, and are weighted to account for selective decisions by the seller to accept an offer.

treatments, buyers are less likely to raise their offers as their own prior offer increases, and the effects are more statistically significant in the B1L and B1H sessions. This may be driven by the sharp reduction in some buyers' offers following seller rejections of offers exceeding K_H, as illustrated in Figure 6.

7 Conclusion

Managers often negotiate in complex and uncertain environments where noisy information flows gradually. This dynamic can alter the degree of information asymmetry between negotiating parties and affect both the efficiency and distribution of bargaining outcomes. Building on the bilateral bargaining model developed by Daley and Green (2020), which features a static lemon condition, we use laboratory experiments to examine how subjects adapt their behavior as news incrementally reveals the seller's private information. In a companion paper, Ding and Lehrer (2025) consider the effects of rivalry among buyers and analyze how competition interacts with subjects' learning behavior in similar laboratory settings.

Our laboratory results show that subjects display both strategic postures and fairness considerations within this dynamic information environment. Anticipating that new information will arrive and gradually reveal the seller's type, buyers hesitate to reach a deal immediately by making conservative initial offers and significantly increase their offers only when enough positive news accumulates. On the other hand, informed sellers are less sensitive to news but primarily adjust their decisions in response to buyers' strategies. Low type sellers display an understanding of and willingness to take advantage of asymmetric information, even when such behavior is in opposition to strong norms of equality. However, as news gradually reduces information asymmetries, low type sellers realize that their chance of being mischaracterized as a high type diminishes, making them more willing to accept offers that equally divide the trading surplus. In contrast, high type sellers benefit less from the information provision. Even as their type is gradually revealed, they continue to face difficulties in securing higher offers that they deem acceptable.

Interestingly, we find that the introduction of information does not increase efficiency, in part because buyers in the control sessions without additional information do not commit to waiting as a screening strategy, resulting in more deals being reached than predicted by theory. However, buyers' aggressive offers often lead to losses when transacting with low type sellers. This result is consistent with prior literature showing that subjects struggle to understand the adverse selection problem even after gaining experience, suggesting a fundamental challenge in contingency reasoning under uncertainty. Nonetheless, such challenges appear to be mitigated when additional information is provided to reduce informational asymmetries.

In summary, as both asymmetric and noisy information flows are ubiquitous in real-world scenarios, it is crucial to understand how individuals adapt their strategies within such environments. Rather than following rational Bayesian updating, decision makers often exhibit cognitive predispositions that lead to biased interpretation or selective use of information. Future research that incorporates these predispositions and behavioral biases into models of dynamic information environments could offer valuable insights for enhancing managerial decision-making, especially if it accounts for behavioral inertia, asymmetric responses to positive and negative news, and fairness considerations documented in our study.

Appendix: Proof of the Model with Inequality Aversion

As discussed in the main text, the equilibrium of our proposed extension has the same structure as Daley and Green (2020). That is, there exists a belief threshold β^γ such that: when $Z_t \geq \beta^\gamma$, the buyer offers $\overline{W}^\gamma$ and both types of seller accept this offer. When $Z_t < \beta^\gamma$, the buyer makes an offer $W_t < \overline{W}^\gamma$ and only an L type seller would be willing to accept such an offer with positive probability. The buyer's value function has a similar form with equation (14) in Daley and Green (2020).

$$F_B(z) = \begin{cases} \big(1-p(z)\big)\left[(1+\gamma)V_L - (1+2\gamma)\overline{W}^\gamma\right] + p(z)(V_H - \overline{W}^\gamma) & \text{if } z \geq \beta^\gamma \\ C_1 \times \frac{e^{u_1 z}}{1+e^z} & \text{if } z < \beta^\gamma \end{cases} \tag{3}$$

The optimal threshold β^γ maximizes

$$F(z|\beta^\gamma) = (1+e^{\beta^\gamma})\Big[\big(1-p(\beta^\gamma)\big)\big((1+\gamma)V_L - (1+2\gamma)\overline{W}^\gamma\big) + p(\beta^\gamma)(V_H - \overline{W}^\gamma)\Big]e^{-u_1\beta^\gamma}\frac{e^{u_1 z}}{1+e^z}$$

at every z. Therefore, $\frac{d}{d\beta^\gamma}F(z|\beta^\gamma) = 0$. Intuitively, β^γ is greater than β in Daley and Green (2020), because a deal reached with an L type seller would incur greater loss.

The parameter C_1 is pinned down by the observation that $F_B(\beta^\gamma) = C_1\frac{e^{u_1\beta^\gamma}}{1+e^{\beta^\gamma}}$. Following Lemma 1 of Daley and Green (2020), the zero net benefit of screening in equilibrium allows us to express

$$(1+2\gamma)W(z) = (1+\gamma)V_L - F_B(z) + \frac{F_B'(z)}{1-p(z)}$$

Combining this expression with equation (3), we can derive that

$$F_L(z) = W(z) = \begin{cases} \overline{W}^\gamma & \text{if } z \geq \beta^\gamma \\ \frac{1+\gamma}{1+2\gamma}V_L + \frac{C_1(u_1-1)}{1+2\gamma}e^{uz} & \text{if } z < \beta^\gamma \end{cases}. \tag{4}$$

Since $\frac{1+\gamma}{1+2\gamma} < 1$, there is some chance that $W(z) < V_L$. In other words, buyers may make some non-experimentation offers, and the probability of employing an experimentation offer decreases in γ. Therefore, this model may help explain and rationalize why in the laboratory we observe very few buyers' offers consistent with experimentation.

Moreover, the acceptance rate at state z is given by

$$q(z) = \frac{rF_L(z) + \frac{\phi^2}{2}F_L'(z) - \frac{\phi^2}{2}F_L''(z)}{F_L'(z)} = (1+\gamma)\frac{\phi^2 V_L}{2C_1}e^{-u_1 z}.$$

The introduction of inequality aversion increases the acceptance rate for L type sellers.

References

Anderson, Steven T., Daniel Friedman, and Ryan Oprea. 2010. "Preemption Games: Theory and Experiment." *American Economic Review*, 100(4): 1778–1803.

Babcock, Linda, George Loewenstein, and Xianghong Wang. 1995. "The Relationship Between Uncertainty, the Contract Zone, and Efficiency in a Bargaining Experiment." *Journal of Economic Behavior & Organization*, 27(3): 475–485.

Barbosa, Luciana, Cláudia Nunes, Artur Rodrigues, and Alberto Sardinha. 2020. "Feed-in Tariff Contract Schemes and Regulatory Uncertainty." *European Journal of Operational Research*, 287(1): 331–347.

Benjamin, Daniel J. 2019. "Errors in Probabilistic Reasoning and Judgment Biases." *Handbook of Behavioral Economics: Applications and Foundations 2 (edited by Doug Bernheim, Stefano DellaVigna, and David Laibson). Elsevier Press.*

Bhagwat, Vineet, Robert Dam, and Jarrad Harford. 2016. "The Real Effects of Uncertainty on Merger Activity." *The Review of Financial Studies*, 29(11): 3000–3034.

Bochet, Olivier, and Simon Siegenthaler. 2018. "Better Later than Never? An Experiment on Bargaining Under Adverse Selection." *International Economic Review*, 59(2): 947–971.

Bochet, Olivier, Manshu Khanna, and Simon Siegenthaler. 2024. "Beyond Dividing the Pie: Multi-Issue Bargaining in the Laboratory." *Review of Economic Studies*, 91(1): 163–191.

Brogaard, Jonathan, Thanh Huong Nguyen, Talis Putnins, and Eliza Wu. 2021. "What Moves Stock Prices? The Role of News, Noise, and Information." *Review of Financial Studies*, 35(9): 4341–4386.

Cason, Timothy N., and Stanley S. Reynolds. 2005. "Bounded Rationality in Laboratory Bargaining with Asymmetric information." *Economic Theory*, 25: 553–574.

Chaigneau, Pierre, Alex Edmans, and Daniel Gottlieb. 2025. "A Theory of Fair CEO Pay." *American Economic Review: Insights*, 7 (3): 306–324.

Coutts, Alexander. 2019. "Good News and Bad News Are Still News: Experimental Evidence on Belief Updating." *Experimental Economics*, 22: 369–395.

Croson, Rachel T. A. 1996. "Information in Ultimatum Games: An Experimental Study." *Journal of Economic Behavior & Organization*, 30(2): 197–212.

Dal Bó, Pedro, and Guillaume R. Fréchette. 2018. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey." *Journal of Economic Literature*, 56(1): 60–114.

Daley, Brendan, and Brett Green. 2020. "Bargaining and News." *American Economic Review*, 110(2): 428–474.

Deneckere, Raymond, and Meng-Yu Liang. 2006. "Bargaining with Interdependent Values." *Econometrica*, 74(5): 1309–1364.

Descamps, Ambroise, Sébastien Massoni, and Lionel Page. 2022. "Learning to Hesitate." *Experimental Economics*, 25: 359–383.

Ding, Tingting, and Steven F. Lehrer. 2025. "Does Competition Impede Learning from News in the Market for Lemons? Evidence from the Lab." *Working Paper, James Madison University.*

Eil, David, and Justin M. Rao. 2011. "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself." *American Economic Journal: Microeconomics*, 3(2): 114–138.

Embrey, Matthew, Guillaume R. Frechette, and Steven F. Lehrer. 2015. "Bargaining and Reputation: An Experiment on Bargaining in the Presence of Behavioural Types." *Review of Economic Studies*, 82: 608–631.

Enke, Benjamin, and Florian Zimmermann. 2019. "Correlation Neglect in Belief Formation." *Review of Economic Studies*, 86(1): 313–332.

Fanning, Jack, and Andrew Kloosterman. 2022. "A Simple Experimental Test of the Coase Conjecture: Fairness in Dynamic Bargaining." *The Rand Journal of Economics*, 51(1): 138–165.

Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory Of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics*, 114: 817–868.

Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10(2): 171–178.

Fréchette, Guillaume, and Sevgi Yuksel. 2017. "Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination." *Experimental Economics*, 20: 279–308.

Goller, Daniel, and Maximilian Späth. 2024. "'Good Job!' The Impact of Positive and Negative Feedback on Performance." *Sports Economics Review*, 8: 100045.

Grosskopf, Brit, Lucas Rentschler, and Rajiv Sarin. 2018. "An Experiment on First-Price Common-Value Auctions with Asymmetric Information Structures: The Blessed Winner." *Games and Economic Behavior*, 109: 40–64.

Haggard, K. Stephen, John S. Howe, and Andrew A. Lynch. 2015. "Do Baths Muddy the Waters or Clear the Air?" *Journal of Accounting and Economics*, 59(1): 105–117.

Huang, Jennie, Judd B. Kessler, and Muriel Niederle. 2024. "Fairness Concerns Are Less Relevant When Agents Are Less Informed." *Experimental Economics*, 27(1): 155–174.

Kagel, John H., and Dan Levin. 1986. "The Winner's Curse and Public Information in Common Value Auctions." *American Economic Review*, 76(5): 894–920.

Kalay, Alon. 2014. "International Payout Policy, Information Asymmetry, and Agency Costs." *Journal of Accounting Research*, 52(2): 457–472.

Keniston, Daniel, Bradley J. Larsen, Shengwu Li, J.J. Prescott, Bernardo S. Silveira, and Chuan Yu. 2024. "Splitting the Difference in Incomplete Information Bargaining: Theory and Widespread Evidence from the Field." *International Economic Review*, 65(4): 1911–1939.

Kisgen, Darren J., Jun "QJ" Qian, and Weihong Song. 2009. "Are Fairness Opinions Fair? The Case of Mergers and Acquisitions." *Journal of Financial Economics*, 91(2): 179–207.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa. 2019. "Failures in Contingent Reasoning: The Role of Uncertainty." *American Economic Review*, 190(10): 3437–3474.

Mitzkewitz, Michael, and Rosemarie Nagel. 1993. "Experimental Results on Ultimatum Games with Incomplete Information." *International Journal of Game Theory*, 22: 171–198.

Müobius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat. 2022. "Managing Self-Confidence: Theory and Experimental Evidence." *Management Science*, 68(11): 7793–8514.

Oprea, Ryan, Daniel Friedman, and Steven T. Anderson. 2009. "Learning to Wait: A Laboratory Investigation." *Review of Economic Studies*, 76(3): 1103–1124.

Rapoport, Amnon, Ido Erev, and Rami Zwick. 1995. "An Experimental Study of Buyer-Seller Negotiation with One-Sided Incomplete Information and Time Discounting." *Management Science*, 41(3): 377–394.

Rapoport, Amnon, James A. Sundali, and Darryl A. Seale. 1996. "Ultimatums in Two-Person Bargaining with One-Sided Uncertainty: Demand Games." *Journal of Economic Behavior & Organization*, 30(2): 173–196.

Reynolds, Stanley S. 2000. "Durable-Goods Monopoly: Laboratory Market and Bargaining Experiments." *The RAND Journal of Economics*, 31(2): 375–394.

Roth, Alvin E. 1995. "Bargaining Experiments." *Handbook of Experimental Economics*, 253–348.

Sandria, Serena, Christian Schade, Oliver Mußhoffb, and Martin Odening. 2010. "Holding on for Too Long? An Experimental Study on Inertia in Entrepreneurs' and Non-Entrepreneurs' Disinvestment Choices." *Journal of Economic Behavior & Organization*, 76: 30–44.

Straub, Paul G., and J. Keith Murnighan. 1995. "An Experimental Investigation of Ultimatum Games: Information, Fairness, Expectations, and Lowest Acceptable Offers." *Journal of Economic Behavior & Organization*, 27(3): 345–364.

Vespa, Emanuel, and Alistair J. Wilson. 2019. "Experimenting with the Transition Rule in Dynamic Games." *Quantitative Economics*, 10(4): 1825–1849.

Weber, Matthias, John Duffy, and Arthur J. H. C. Schram. 2018. "An Experimental Study of Bond Market Pricing." *The Journal of Finance*, 73(4): 1857–1892.

Wittenberg-Moerman, Regina. 2008. "The Role of Information Asymmetry and Financial Reporting Quality in Debt Trading: Evidence from the Secondary Loan Market." *Journal of Accounting and Economics*, 46(2): 240–260.

Online Appendix for

Waiting for the Right Offer: Laboratory Evidence on How News Affects Bargaining

Tingting Ding
James Madison University

Steven F. Lehrer
Queen's University and NBER

CONTENTS:

A Experimental Instructions

A.1 Instructions for Treatment C1

You are about to participate in a session on decision-making, and you will be paid for your participation in cash vouchers, privately at the end of the session. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

Please turn off your cellular phone(s) now. Please close any programs that you may have opened on the computer. The entire session will take place through computer terminals, and all interaction between you and other session participants will take place through the computers. Please do not talk directly to or attempt to communicate with other participants during the session.

We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the session and will be shown how to use the computers. If you have any questions during this period, raise your hand and your question will be answered in private.

Instructions

In this experiment you will be asked to make decisions in 12 trading periods. At the beginning of the experiment you will be randomly assigned to be a buyer or a seller and remain in this same role in each trading period.

At the beginning of each trading period subjects are matched at random together to form pairs, with one buyer and one seller in each pair. Each buyer will have an opportunity to make one purchase of a fictitious perishable commodity from the seller they are matched with during the period, according to trading rules that are explained below.

In each trading period, a fictitious perishable commodity will be exchanged. There is a 2/3 chance that the commodity is type Green and a 1/3 chance that the commodity is type Blue. Only the seller knows whether the commodity is type Green or type Blue. Buyers do not observe the color type of the good.

A buyer will make an offer for the fictitious commodity at the start of each period. Each buyer values a type Blue good at 300 and values a type Green good at 60.

Once the buyer makes an offer, the seller will choose whether to accept or reject this offer. Each seller faces an initial cost of selling a good that differs by the color type of good. A good of type Blue has an initial cost of 180 and a good of type Green has an initial cost of 0.

If the seller accepts the initial offer of the buyer, a trade is made and the period ends.

If the seller rejects the initial offer of the buyer, the buyer-seller matched pair moves to the second stage which may consist of multiple rounds. The attached figures show the computer screen in Stage 2. In each round of the second stage, the buyer will have 6 seconds to make an offer in that round. Buyers make decisions by moving their mouse along the slider bar in the middle of the screen. At the start of each round, the mouse is located at the position of the last offer. By moving along the slider bar to a desired location, the buyer can make a new offer. The amount of this new offer is displayed on the screen and by hitting the OK button, this offer is made. However, if the buyer did not hit the OK button within 6 seconds, the computer will treat the final position of the slider as the desired offer. If the buyer does not move the slider bar within the round, the computer will assume the buyer wishes to make the exact same offer that he made in the previous round. Thus, after 6 seconds, the computer will record a new offer from the buyer. The seller will then have 6 seconds to make a decision on the buyer's updated decision. As before, if the seller rejects the updated offer, the matched pair would move to another round in the second stage. Only if the seller accepts the offer, a trade is made and the period ends. In the event of the seller not reaching a decision within 6 seconds, the computer will make a default decision of rejection in that round of the second stage.

We will soon provide some numerical examples to clarify how much the buyer and seller earn if a trade is made. We will also conduct 4 practice periods to familiarize you with the experiment as seen by both the buyer and the seller.

The opportunity to reach a deal in the second stage might evaporate before you seize it, in which case both the buyer and seller earn 0 points that period. This can happen since at the end of each round there is a 4% chance that the period will end. The computer will draw a random number from 0 to 100, where each number is equally likely to have been drawn. If the computer draws a number larger than 96, this is the round determined to have ended the period.

If you enter the second stage, you will play in blocks of 12 rounds. At the end of each block of 12 rounds, the computer will report if a number greater than 96 was drawn in a prior round in that block and if so, in which round. In this case, that round ends the period, irrespective of whether a deal was subsequently reached in a later round. If the computer reports that none of the twelve numbers drawn in the block were larger than 96, the players in the group will move to a next block of 12 second stage rounds. In other words, the total number of rounds in the second stage is determined randomly by the computer.

Once the period ends, the color type of the good being traded is announced to the pair. If a deal is reached, the profit or loss of a buyer is calculated as the difference between his value and the offer he made. The profit or loss of a seller is calculated as the difference between the amount she received and her cost.

A few examples might help your understanding. These are not meant to be realistic:

1. At the start of the period, the seller is informed she is selling a Blue type good. The buyer makes an initial offer of 215 that the seller immediately accepts. The period ends immediately. Since this is a Blue good, the buyer receives 300-215=85 points and the seller receives 215-180=35 points. If instead the seller had a Green type of good and immediately accepted the offer of 215, the buyer would receive 60-215= -155 points and the seller would receive 215-0=215 points.

2. At the start of the period, the seller is informed she is selling a Blue type good. The buyer makes an initial offer of 160 that the seller rejects. The matched pair moves to the second stage. After 12 rounds, the buyer makes a revised offer of 220 that the seller accepts. In this case, if none of the random numbers drawn in the prior rounds was greater than 96, the deal stands, the buyer earns 300-220=80 points and the seller earns 220-180=40 points. However, if the computer randomly drew a number greater than 96 in one of the prior rounds, say if the computer drew 98 in round 10, the deal in round 12 is canceled, in which case both the buyer and seller receive 0 points. If instead, the seller does not accept any offers in round 12, there is a chance the matched pair moves to another block of twelve rounds. This chance happens if the computer did not draw a number larger than 96 in any of the earlier rounds in the block.

3. At the start of the period, the seller is informed she is selling a Green type good. The buyer makes an initial offer of 60 that the seller rejects. The matched pair moves to the second stage. After 10 rounds, the buyer makes an offer of 85 that the seller accepts. In this case, if none of the random numbers drawn in the prior rounds was greater than 96, the deal stands, and the buyer earns 60-85=-25 points and the seller earns 85-0=85 points. However, if ever a random number drawn in the prior rounds is greater than 96, say if the computer drew 98 in round 3, the deal in round 10 is canceled, in which case both the buyer and seller in this group receive 0 points. If instead, after 10 rounds, the seller rejects this revised offer of 85, the matched pair moves to round 11.

As you can see there are many possibilities.

When every matched pair has finished this task, and completed a transaction for the fictitious good, the next period begins. You will be randomly re-assigned to a player in the next period. You will not be able to identify whom you have paired with in previous or future periods. You will remain in the same role of either a buyer or seller in all periods and the task in every period is exactly the same as the one just described (but with the randomly re-matched player). The color type of the commodity will be randomly assigned in each period. That is, regardless of what type of commodity a seller holds previously, in the next period the seller still has a 1/3 chance of having a Blue type of commodity and a 2/3 chance of having a Green type of commodity. The session consists of 12 such periods.

At the end of the experiment you will be called individually and paid in cash voucher in private by the experimenter. The amount of your cash payment includes a $10 participation fee and your earnings in the experiment. Your earning in the experiment will be determined as follows: In the beginning of each period in today's session each buyer and seller will receive an endowment of 120 points. At the end of the session, one of the 12 trading periods will be randomly selected, and you will earn the total number of points that you obtained from the trade in that period plus your initial endowment. The conversion rate from points to U.S. dollars is $1 = 12 points. For instance, if in the chosen period you are a buyer and earn 48 points from the trade, you will be paid a $10 participation fee and (48 points + 120 points) /12 = $14 from the trade. In total, you will be paid $24. However, if your loss exceeds 120 points in the randomly selected period, you will only be paid the $10 participation fee.

Are there any questions?

Summary

Before we start, let me remind you that:

- After a period is finished, you will be randomly re-matched to a player of the opposite role for the next period.

- In each period, you and another player will form a pair to complete a transaction of a fictitious good. The color type of commodity will be randomly assigned in each period. Only the seller knows the color type of the good at the beginning of the period.

- If the seller accepts the buyer's initial offer, the period ends. If the seller rejects the buyer's initial offer, you move to a second stage. In each round of the second stage, buyers can decide whether to change their offer by moving along the slider bar and sellers can decide whether or not to accept or reject an updated offer. You have 6 seconds to act. Each round lasts 12 seconds.

- Each round of the second stage ends (with or without a deal being reached) with a small probability of 4% that is announced after a block of 12 rounds.

- At the end of the session, your earnings are determined by the number of points you earned in a randomly selected period plus a starting balance that depends on the role you play. In addition, you will receive a $10 participation fee.

Good Luck.

Figure A.1: Screenshot of Buyer in the C1 Treatment

Figure A.2: Screenshot of Seller in the C1 Treatment

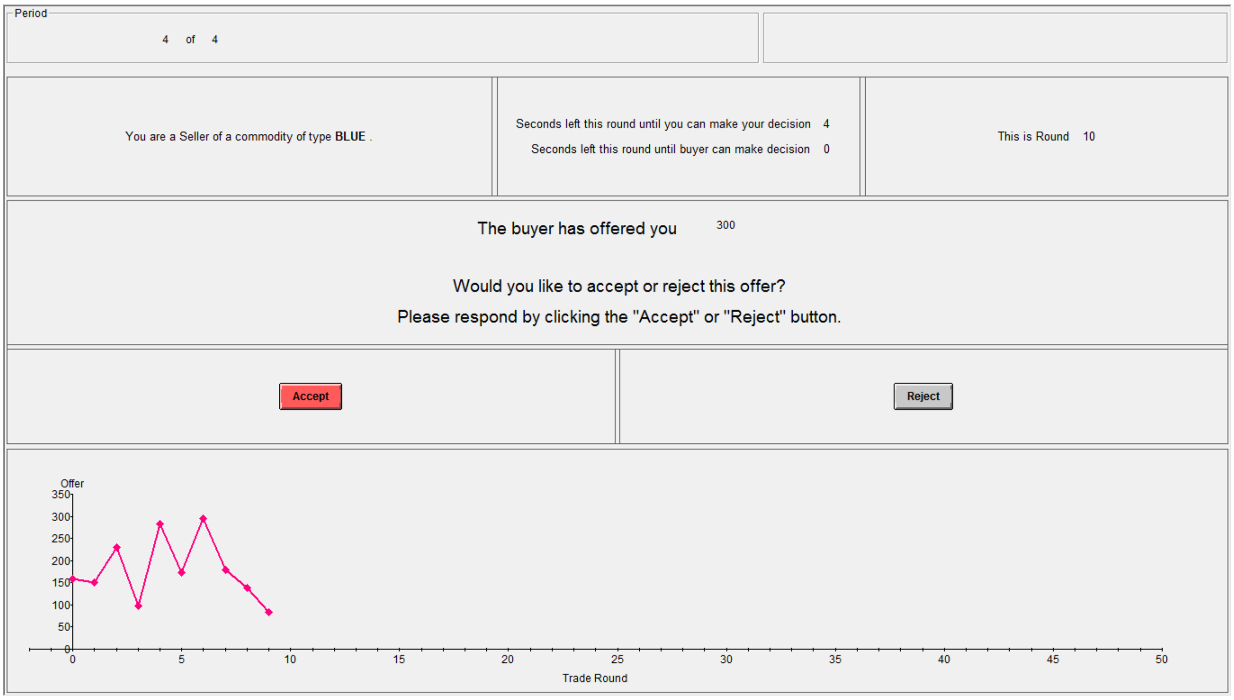

A.2 Instructions for Treatment B1L

. . . . . .

We will soon provide some numerical examples to clarify how much the buyer and seller earn if a trade is made.

In each round of the second stage, for each matched pair, the bottom panel of the screen will display the information about the color of the good the seller is offering and the past offers made. While the seller knows the color type of the good, this information may help the buyer learn about the color type. To understand how information is presented, the key features are:

- Information starts at 0.
- Each second stage round one new tick is displayed.
- The tick either moves up or down by a fixed amount, 0.01. That is, the information either takes a step up or a step down from its prior value.
- When the good is type Green, the tick is more likely to move down than up. On average, it moves down with probability 55%, and moves up with probability 45%.
- When the good is type Blue, the tick is more likely to move up than down. More precisely, it moves up with probability 55%, and moves down with probability 45%.

In the practice periods, several examples of how information is presented in the second stage rounds for both Blue and Green goods will be shown. We will also conduct 4 practice periods to familiarize you with the experiment as seen by both the buyer and the seller.

. . . . . .

A.3 News Displayed in the Lab

Figures A.3 and A.4 respectively present the pre-generated news in the treatments with information. The blue lines indicate the news gradually revealed to H type sellers, while the green lines indicate the news gradually revealed to L type sellers.

Figure A.3: News Displayed in the B1H Treatment

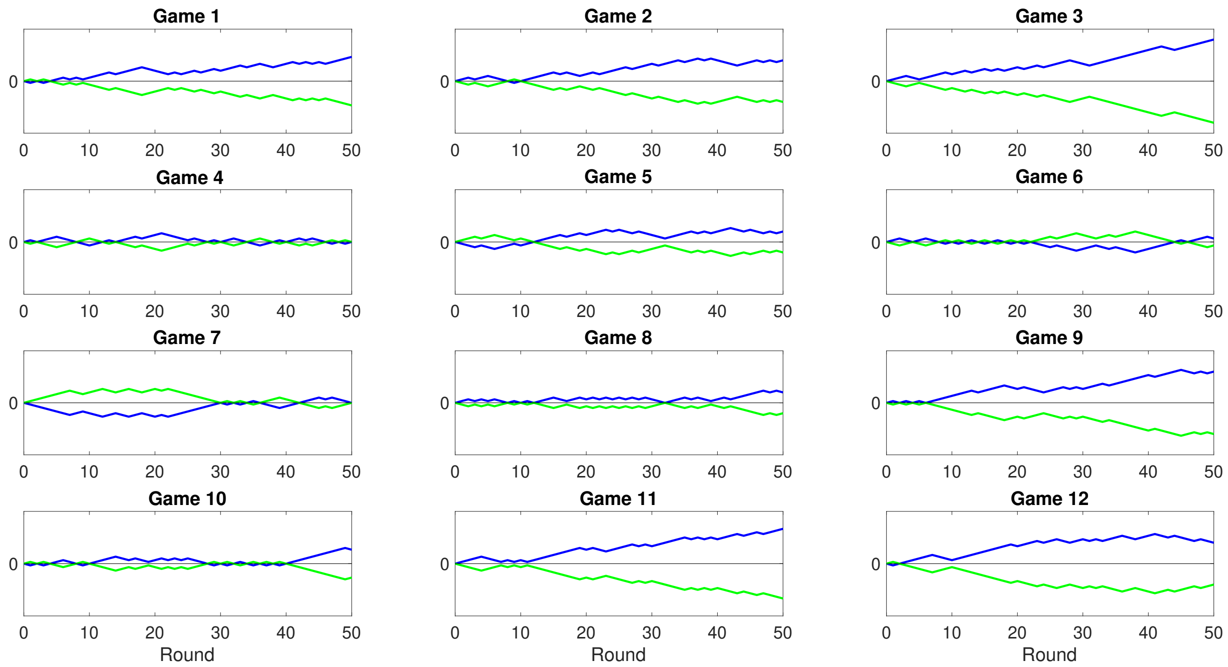

Figure A.4: News Displayed in the B1L Treatment

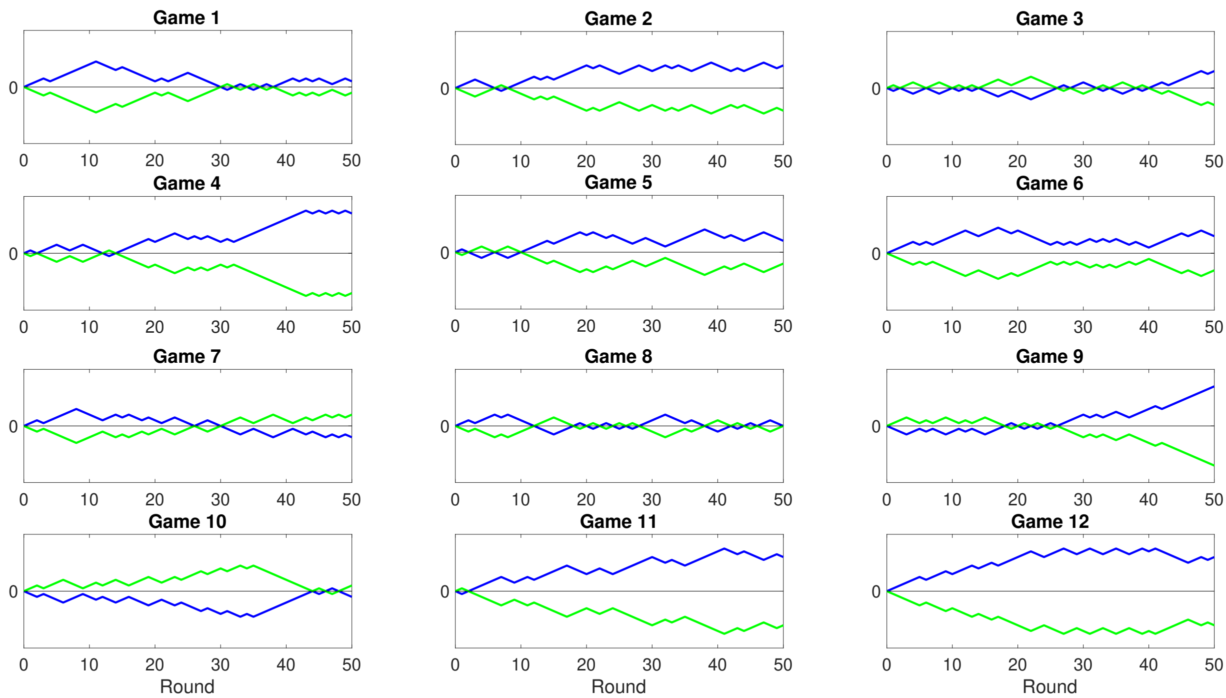

B Additional Empirical Results

In this subsection, we summarize additional results that provide further insights into how sellers and buyers make decisions within the experimental environment. As noted in the main text, we first consider an alternative specification that explains how the buyer's offer, W_{irg}, evolves between successive rounds:

$$\Delta W_{irg} = \alpha_1 Up_{irg} + \alpha_2 News_{irg} + \alpha_3 r_r + \alpha_4 I_i + \varepsilon_{irg} \tag{B.1}$$

This specification differs from equation (**??**) in that we exclude the prior offers as an explanatory variable, due to econometric concerns related to the inclusion of a lagged dependent variable as an exogenous regressor.

Table B.1 presents the estimates from equation (B.1), reweighted for selective continuation. Compared to the results reported in Table 6, we find that when prior offers are dropped from the set of explanatory variables, buyers in both the B1L and B1H sessions significantly increase their offers in response to the most recent uptick and the stock of news. This is likely because lagged offers are correlated with other explanatory factors, including prior upticks contained in the News variable, and may interact with recent signals in complex ways.

We next examine how alternative streams of news affect the change in buyers' offers by considering an expanded empirical model. To motivate this strategy, we begin by assuming that the offer made in any round r can be modelled as

$$W_{irg} = h_r(Up_{i1g} ... Up_{irg}, News_{irg}, I_{i0g} ... I_{irg}, v_i, g_g, \varepsilon_{i0g} .. \varepsilon_{irg}), \tag{B.2}$$

where h_r is an unknown twice-differentiable function, $(Up_{i1g} ... Up_{irg})$ are the round specific indicators if the tick moved up, $News_{irg}$ is the indicator for the level of the news stock at the start of the round r, $(I_{i0g} ... I_{irg})$ is the full history of individual characteristics, and $(\varepsilon_{i0g} .. \varepsilon_{irg})$ are independent random shocks. Note that v_i is included to capture unobserved time-invariant individual attributes and g_g is capturing round-invariant game attributes that may in part reflect aggregate subject learning. In our setting, the individual characteristics include gender and major, allowing us to suppress the g and r subscript.

To ease exposition, we consider a setting consisting of just two stage 2 rounds and linearize the offer process in each round. We omit the cumulative stock of news, since its inclusion is statistically insignificant once we control for the last two recent news and their interaction, thereby allowing us to represent the offers in round 1 as

$$W_{i1g} = v_i + \beta_1 Up_{i1g} + \beta_2 I_i + g_g + \varepsilon_{i1g}, \tag{B.3}$$

and the offer in round 2 as

$$W_{i2g} = v_i + \alpha_1 Up_{i1g} + \alpha_2 Up_{i2g} + \alpha_3(Up_{i1g} * Up_{i2g}) + \alpha_4 I_i + g_g + \varepsilon_{i2g}. \tag{B.4}$$

In equation (B.4), we allow for potential complementarities in the news process since news in the first round can interact in unknown ways with the news received at the start of the second round. Notice the notation for the coefficients in equations (B.3) and (B.4) are different, as we do not restrict the effects of observed inputs and news on offers to be constant across rounds.

The above formulation is a triangular model structure. Since all of the explanatory variables are discrete dummy variables, the only restrictive assumption imposed by linearization of equation (B.2) is additive separability of the error terms. Thus, full information maximum likelihood parameter estimates of equations (B.3) and (B.4) are equivalent to equation-by-equation OLS, which does not impose any assumptions on the distribution of the residuals. Ding and Lehrer (2010) point out that a local insensitivity assumption introduced in Chesher (2003) is needed to nonparametrically identify all the structural parameters in this triangular system of offer equations in both levels and first differences to remove v_i.

This formulation is easy to extend beyond two rounds but could lead to a large number of structural parameters if there are complementarities between news across all or any combination of rounds. In Table B.2 we consider t-tests from specifications that include further lags of the news process, and the results suggest that offers significantly change in response to only the two most recent pieces of news. This allows us to exclude news released in earlier rounds in the remaining specifications.[1] To estimate the heterogeneous effects of Up signals as well as potential dynamic complementarities between the most recent Up signals, we represent the data in each round r in first differenced format as

$$\Delta W_{irg} = (\alpha_1^r - \alpha_2^{r-1})Up_{i(r-1)g} + \alpha_2^r Up_{irg} + \alpha_3^r(Up_{i(r-1)g} * Up_{irg}) + (\alpha_4^r - \alpha_4^{r-1})I_i + t_r + \varepsilon_{irg}^* \tag{B.5}$$

where $\varepsilon_{irg}^* = \varepsilon_{irg} - \varepsilon_{i(r-1)g}$. The parameter t_r can be viewed as capturing common differences in round effects and as a sufficient statistic of where the tick stands in that round, since we ensured that news followed the same process in the same game across sessions.

As in the main text, we continue to use inverse probability weighting to recover unbiased and consistent parameter estimates of equation (B.5). This analysis mimics that undertaken in the main text but controls for a different set of variables in the news

[1]Similarly, F-tests from richer specifications almost always reject the inclusion of additional lagged up ticks as well as further complementarities.

process. We also focus on data from the first block of 12 rounds in a game since data from subsequent blocks need to be reweighted. Table B.3 presents estimates of equation (B.5) in which we restrict the coefficients to be constant across rounds. Each column presents results corresponding to a different treatment. We find that buyers significantly increase their offers when the most recent signal is up (in the B1H sessions) or when they receive two straight up signals (in the B1L sessions).[2]

Further, we attempt to shed additional light on how the up ticks in the news process influence the change in offers between successive rounds, since the control group in the above test consists of individuals who either received just one or zero up ticks in the two most recent rounds. We consider the following specification,

$$\begin{aligned}\Delta W_{irg} =& \gamma_1(I[Up_{irg}=1]*I[Up_{i(r-1)g}=0])+\gamma_2(I[Up_{irg}=0]*I[Up_{i(r-1)g}=0])\\ &+\gamma_3(Up_{i(r-1)g}*Up_{ig})+\gamma_4 I_i+\varepsilon_{irg}\end{aligned} \tag{B.6}$$

where the first term contains indicator functions to capture subjects who received an up tick in the current round, but in the prior round received a down tick. A special case of this specification would involve estimation of a contemporaneous model in first differences as expressed by

$$\Delta W_{irg} = \delta_1 \Delta Up_{irg} + \delta_4 I_i + \varepsilon_{irg} \tag{B.7}$$

The first differenced model would impose restrictions on $\Delta W_{irg} = \phi_1(I[Up_{irg} = 1] * I[Up_{i(r-1)g} = 0]) + \phi_2(I[Up_{irg} = 0] * I[Up_{i(r-1)g} = 0] + \phi_3(Up_{i,(r-1)g} * Up_{ig}) + \phi_4 I_i + \varepsilon_{irg}$ That is, there is only a response to the current news ($\phi_3 = 0$) and there are offsetting effects from any change in news ($\phi_1 = -\phi_2$). Similarly, a fixed effects specification of a buyer offer equation would also impose an assumption of equal offsetting effects that our main analysis relaxed.

Table B.4 presents estimates of equation (B.6) as well as the results from specification tests. The specification tests consistently reject the null hypothesis that would support a restricted first differenced model. In addition, Table B.4 indicates that buyers underreact to the news, as they adjust their offers by a smaller magnitude than theoretical predictions. Compared to the scenario of receiving two down ticks, Daley and Green (2020)

[2]Our reduced form analysis does share some features with Enke et al. (2023) who collect data on beliefs and regress them on the current news signal as well as the prior signal. However, we only have information on actions (buyers' offers and seller decisions) but not their beliefs. This is a result of how we implement the theoretical model in discrete time, whereby we allow subjects to have sufficient time to process the information. Collecting information on beliefs within this setup would lead the discrete time analog to have a longer reaction time and reduce its connection to the underlying continuous time model. An advantage of our experimental setup is that within each period, analyzing decisions after the third round permits us to disentangle the impact of lagged news signals from the stock value of all prior signals in that round, as well as allowing the consideration of dynamic complementarities in the news signal process.

predicts that buyers in the B1L treatment should increase offers by 29.79 when receiving two up ticks, increase by 20.28 when receiving an up tick at round t and a down tick at round $t-1$, and decrease by 8.32 when receiving a down tick at round t and an up tick at round $t-1$. However, we observe the corresponding magnitudes in these sessions are 10.76, 1.90 and -0.64, respectively. Similar evidence is also found in the B1H treatment.[3]

A further examination of the data is available from the authors upon request. This additional analysis involves robustness checks including different specifications and assumptions on the seller's acceptance decision equation. Linear and nonlinear econometric estimation of reduced form models that explain both the buyer's offer and the seller's acceptance are considered. Briefly, there are four main findings. First, we find that sellers focus on the aggregate stock of news and not the individual up ticks, particularly early within each session. In contrast, buyers focus only on recent up ticks and later in the session react more strongly to dynamic complementarities in the up ticks. Second, the econometric estimates confirm that sellers react differently to the news information process if assigned an H type versus an L type good for sale. Third, robustness exercises confirm that subjects under react to the news in the bilateral bargaining sessions. Fourth, for many round groupings in both the B1L and B1H treatments, DuMouchel and Duncan (1983) tests support the use of sampling weights when estimating equation (B.5) and thereby account for selective seller behavior. This provides further evidence of a statistically significant difference, and both the weighted and unweighted estimates show that buyers and sellers respond to different portions of the news process. The additional analysis also presents results using data from all games.

[3]Compared to the scenario of receiving two down ticks, theory predicts that buyers in the B1H treatment should increase offers by 22.40 when receiving two up ticks, increase by 16.97 when receiving an up tick at round t and a down tick at round $t-1$, and decrease by 5.98 when receiving a down tick at round t and an up tick at round $t-1$. But we observe the corresponding magnitudes in these sessions are 6.96, 3.99 and -0.68, respectively.

Table B.1: Factors Influencing How Buyers Adjust Their Offers Across Rounds

Treatment	**C1**	**B1L**	**B1H**
Up Tick at t		4.32^{***}	4.30^{**}
		(1.37)	(1.72)
News Stock		72.23^{***}	90.82^{***}
		(15.07)	(28.19)
Round	0.02	0.10^{*}	-0.07
	(0.05)	(0.06)	(0.09)
Constant	8.27^{***}	2.73^{**}	4.26^{***}
	(1.39)	(0.94)	(1.15)
Control	Yes	Yes	Yes
Observations	1,745	1,844	1,706
R-squared	0.009	0.021	0.013

Note: Robust standard errors in parentheses clustered at the session level. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$. The models assume a homogeneous effect of each parameter in the offer equation and are weighted to account for selective decisions by the seller to accept an offer.

Table B.2: Specification Tests of Regressions That Additionally Include a Two-Round Lagged Up Tick Indicator Assuming Homogeneous Impacts

Treatment	**B1L**		**B1H**	
	All Rounds	First Block	All Rounds	First Block
t Statistic	0.99	0.16	2.06**	1.45
p-value	0.32	0.87	0.04	0.15

Note: In this table we present results from a t-test of statistical significance on the coefficient of the two-round lagged up tick variable that is the sole addition to the specifications in equation (B.5). The p-value corresponds to a two-sided test using clustered standard errors at the subject level.

Table B.3: Factors Influencing How Buyers Adjust Their Offers When Dynamic Complementarityis Incorporated

Treatment	**C1**	**B1L**	**B1H**
Up Tick		0.978	4.468**
		(1.543)	(1.618)
Up Tick Prior Round		−1.243	−0.423
		(1.984)	(1.299)
Two Straight Up Ticks		11.300***	3.063
		(3.497)	(3.174)
Constant	8.361***	2.562**	3.263***
	(1.733)	(0.959)	(0.408)
Control	Yes	Yes	Yes
Observations	1,745	1,708	1,570
R-squared	0.007	0.024	0.011

Note: Robust standard errors in parentheses clustered at the session level. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$. The models assume a homogeneous effect of each parameter in the offer equation and are weighted to account for selective decisions by the seller to accept an offer.

Table B.4: Regression Results that Examine How
Alternative Sequences of News Influences Changes in Buyer Offers

Treatment	**B1L**		**B1H**	
	All Rounds	First Block	All Rounds	First Block
Two Straight	10.764***	8.844***	6.956***	8.164***
Up Ticks	(1.891)	(2.165)	(2.135)	(2.998)
Up Tick at $t-1$	-0.641	0.915	-0.682	-2.266
Down Tick at t	(2.140)	(2.459)	(2.120)	(2.148)
Down Tick at $t-1$	1.899	1.360	3.985*	4.023
Up Tick at t	(1.668)	(1.954)	(2.091)	(2.625)
Constant	1.851	2.213	2.733	1.156
	(1.600)	(2.141)	(1.690)	(1.604)
Control	Yes	Yes	Yes	Yes
F-test of Restrictions	18.94	8.39	5.41	4.60
on FD model	[0.00]	[0.00]	[0.01]	[0.01]
Observations	1,708	1,169	1,570	1,107
R-squared	0.022	0.018	0.010	0.018

Note: Robust standard errors in parentheses. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$.

C Robustness Analyses

This section provides robustness checks using two alternative samples: (i) all trading games and (ii) all subjects, including those excluded in the main analysis. As detailed below, the results remain qualitatively consistent with those reported in Table 3.

C.1 Results Using All Games

Incorporating all trading games, Table C.1 shows that descriptive statistics remain consistent with the last 6 games analyzed in Table 3. Nevertheless, we observe several tendencies as subjects gain experience:

- Subjects across all three treatments tend to make more generous initial offers and are more likely to adopt experimentation strategies in the early games.
- Buyers in the C1 sessions initially make more aggressive offers during the first 6 games. Although experience helps to mitigate this tendency, they continue to suffer losses when trading with L type sellers in the latter 6 games.
- In the B1L and B1H sessions, H type sellers initially accept lower offers and reach agreements faster in the first 6 games. However, as they gain experience, they appear to learn how to leverage the information to improve their bargaining positions, although this often results in delays in reaching a deal.

Table C.1: Descriptive Statistics Across All Trading Games

Treatment		**C1**	**B1L**	**B1H**
First Offer		36.56	26.71	23.03
		(38.27)	(24.25)	(17.97)
Accepted Offer				
	H Type Seller	206.25	210.25	198.88
		(52.35)	(38.24)	(47.62)
	L Type Seller	110.34	71.27	66.74
		(67.31)	(61.96)	(62.94)
Payoffs				
	Buyer	−11.83	9.78	16.59
		(79.27)	(62.26)	(68.76)
	H Type Seller	15.39	16.38	11.60
		(41.98)	(31.85)	(38.33)
	L Type Seller	97.87	60.88	61.18
		(72.41)	(62.55)	(63.02)
Deal Round				
	H Type Seller	12.84	16.46	16.23
		(9.27)	(10.20)	(9.97)
	L Type Seller	7.88	8.24	6.06
		(7.12)	(6.69)	(5.21)
Fraction of Deals Reached				
	H Type Seller	0.59	0.54	0.61
		(0.50)	(0.50)	(0.49)
	L Type Seller	0.89	0.85	0.92
		(0.32)	(0.35)	(0.28)
Efficiency		42.26	37.10	42.80
		(15.50)	(12.52)	(12.60)
Experimentation Offers		0.18	0.07	0.06
Fraction of offers $\in (V_L, K_H)$		(0.39)	(0.26)	(0.23)

Note: Each entry presents the mean and standard deviation of the row variable for the treatment denoted by the column heading. The variables deal round and accepted offer are recorded for the sample in which deals were reached.

C.2 Results Using All Subjects

This subsection examines the last 6 games, including two previously excluded subjects: a C1 seller who accepted any offer in the first stage or in the first round of Stage 2, and a C1 buyer who consistently offered 190 points. As both subjects were in the C1 sessions, their inclusion does not affect our hypothesis testing. However, we are concerned that their distinct behavioral patterns could misrepresent typical bargaining behavior in the descriptive analysis. For instance, as shown in Table C.2, the excluded buyer's consistent offering of 190 points increases the average first offer from 27.47 to 35.58, while the average payoffs of H type sellers decrease from 5.18 to -4.15, driven by the excluded seller's frequent acceptance of extremely low offers in early rounds.

Table C.2: Descriptive Statistics for the Last 6 Trading Games (All Subjects)

Treatment		**C1**	**B1L**	**B1H**
First Offer		35.58	25.13	18.27
		(43.43)	(21.63)	(10.83)
Accepted Offer				
	H Type Seller	172.04	213.74	208.08
		(70.64)	(14.49)	(25.92)
	L Type Seller	111.69	71.63	60.87
		(71.62)	(63.38)	(63.96)
Payoffs				
	Buyer	-6.86	5.08	15.46
		(89.64)	(56.79)	(60.79)
	H Type Seller	-4.15	13.35	14.63
		(50.64)	(18.93)	(23.33)
	L Type Seller	94.24	58.20	52.63
		(77.34)	(63.60)	(63.00)
Deal Round				
	H Type Seller	11.28	17.16	18.92
		(9.81)	(10.74)	(10.77)
	L Type Seller	9.02	8.38	7.19
		(7.81)	(7.39)	(5.79)
Fraction of Deals Reached				
	H Type Seller	0.52	0.40	0.52
		(0.50)	(0.49)	(0.50)
	L Type Seller	0.84	0.81	0.86
		(0.36)	(0.39)	(0.34)
Efficiency		36.47	32.46	36.84
		(16.01)	(11.00)	(11.23)
Experimentation Offers		0.13	0.06	0.04
Fraction of offers $\in (V_L, K_H)$		(0.34)	(0.23)	(0.19)

Note: Each entry presents the mean and standard deviation of the row variable for the treatment denoted by the column heading. The variables deal round and accepted offer are recorded for the sample in which deals were reached.

C.3 Robustness to Alternative Approaches for Statistical Inference

This subsection addresses issues related to statistical inference, as our main analysis employs cluster-robust standard errors in linear regression models, with inference based on the Student's t-distribution. Motivated by difference-in-differences estimators commonly used in natural experiments, a body of econometric research has shown that, when the number of clusters is small or cluster sizes are unequal, cluster-robust standard errors may be underestimated, leading to over-rejection of the null hypothesis.

To address this concern, critical values based on the wild cluster bootstrap have been shown to perform more reliably when the number of clusters is moderately large, typically around 12 or more. Furthermore, Monte Carlo evidence from Webb (2023) suggests that the wild cluster bootstrap using the Webb six-point distribution performs well even with as few as five clusters. Table C.3 presents the analog to Table 6 in the main text, using this version of the wild cluster bootstrap. We use a large number of bootstrap replications (9,999) to ensure that the approximate test achieves a rejection frequency close to the nominal 5% level. The p-values in square brackets indicate statistically significant effects of the explanatory variables at levels similar to those reported in Table 6.

Table C.4 presents the analog to Table 6 in the main text, using standard errors calculated via the delete-cluster jackknife. Notably, the jackknife standard errors in Table C.4 are approximately 30–50% larger than the cluster-robust standard errors in Table 6. Only the effect of Prior Offer in the C1 treatment is no longer statistically significant. These jackknife standard errors, proposed by Hansen (2024), are designed to mitigate the worst-case downward bias associated with traditional cluster-robust standard errors and to improve confidence interval coverage.

The proposed jackknife variance estimator is formally shown to perform well under arbitrary cluster sizes, within-cluster correlation, and heteroskedasticity, among other complexities. In cases with a small number of clusters, Hansen (2024) recommends reporting a 95% jackknife confidence interval using an adjusted degrees-of-freedom (DoF) approximation. These adjustments aim to reduce size distortions and involve computing both an adjusted DoF and a scale correction. In our analysis, we find that for the B1L treatment, the adjusted degrees of freedom range from 1.81 to 2.86, while the scale adjustments range from 1.17 to 1.38 across explanatory variables. In general, these adjustment factors were modest across all three treatments, which explains why the main conclusions remain consistent even when using conventional confidence intervals based on jackknife standard errors. As noted earlier, most coefficients that were statistically significant at or below the 5% level in Table 6 remain so with the adjusted p-values in Table C.4.

We speculate that the few observed differences across inference methods might not be surprising since Monte Carlo studies on unbalanced cluster sizes have shown that over-rejection tends to occur when one or a few clusters are unusually large. In our case, the B1L treatment has cluster sizes ranging from 358 to 560, with a mean of 461, while the B1H treatment ranges from 350 to 492, with a mean of 426.5. These ranges indicate no extreme disparities in session sizes within or across treatments. Moreover, we believe the limited differences in inference results across methods are partly due to the experimental design, in which many explanatory variables are randomly assigned in each round. This leads to relatively low within-cluster correlation compared to typical observational datasets. Taken together, these features increase our confidence in the robustness of the findings reported in the main text.

Table C.3: Robustness of Factors Influencing Buyer Offer Adjustments Across Rounds: Wild Cluster Bootstrap P-values

Treatment	**C1**	**B1L**	**B1H**
Up Tick		2.16	3.28**
		(0.85)	(4.90)
		[0.42]	[0.046]
News Stock		155.52**	207.00**
		(4.51)	(3.62)
		[0.02]	[0.03]
Prior Offer	−0.12*	−0.17**	−0.16**
	(2.88)	(−4.59)	(−5.26)
	[0.05]	[0.02]	[0.049]
Round	0.17*	0.47**	0.22
	(1.65)	(6.60)	(0.95)
	[0.05]	[0.04]	[0.43]
Constant	14.04**	7.13	8.75**
	(16.55)	(2.70)	(5.49)
	[0.02]	[0.18]	[0.04]
Control	Yes	Yes	Yes
Observations	1,745	1,844	1,706
R-squared	0.037	0.071	0.051

Note: P-values calculated using the wild cluster bootstrap with Webb weights are in square parentheses and are calculated using 9,999 bootstrap replications. Wild bootstrap-t, null imposed are reported with 3 degrees of freedom in regular parentheses. The models include the full set of subject demographic characteristics and assume a homogeneous effect of each parameter in the offer equation.

Table C.4: Robustness of Factors Influencing Buyer Offer Adjustments Across Rounds: Jackknife Standard Errors and Adjusted p-Values

Treatment	**C1**	**B1L**	**B1H**
Up Tick		2.16	3.28**
		(3.23)	(1.05)
		[0.49]	[0.04]
		{−6.80, 11.22}	{0.37, 6.18}
News Stock		155.52**	207.00*
		(50.01)	(84.52)
		[0.04]	[0.07]
		{16.87, 294.17}	{−30.47, 444.47}
Prior Offer	−0.12	−0.17**	−0.16**
	(0.06)	(0.05)	(0.05)
	[0.11]	[0.03]	[0.047]
	{−0.31, 0.06}	{−0.30, −0.03}	{−0.31, −0.01}
Round	0.17	0.47**	0.22
	(0.26)	(0.13)	(0.35)
	[0.46]	[0.04]	[0.47]
	{−1.05, 1.39}	{0.04, 0.10}	{−0.80, 1.25}
Constant	14.04***	7.13	8.75**
	(2.15)	(3.52)	(2.10)
	[0.005]	[0.10]	[0.02]
	{8.50, 19.59}	{−2.96, 17.22}	{2.86, 14.64}
Control	Yes	Yes	Yes
Observations	1,745	1,844	1,706
R-squared	0.037	0.071	0.051

Note: Hansen (2024) delete-cluster jackknife standard errors are in parentheses, where the clustering is at the session level. In square parentheses, we present p-values calculated using the adjusted Student t distribution; in braces we present 95% confidence intervals calculated using the adjusted Student t distribution. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$. The models include the full set of subject demographic characteristics and assume a homogeneous effect of each parameter in the offer equation.

References

Chesher, Andrew. 2003. "Identification in Nonseparable Models." *Econometrica*, 71(5): 1405–1441.

Daley, Brendan, and Brett Green. 2020. "Bargaining and News." *American Economic Review*, 110(2): 428–474.

Ding, Weili, and Steven F. Lehrer. 2010. "Estimating Treatment Effects from Contaminated Multi-Period Education Experiments: The Dynamic Impacts of Class Size Reductions." *Review of Economics and Statistics*, 92(1): 31–42.

DuMouchel, William H., and Greg J. Duncan. 1983. "Using Sample Survey Weights in Multiple Regression Analyses of Stratified Samples." *Journal of the American Statistical Association*, 78(383): 535–543.

Enke, Benjamin, Uri Gneezy, Brian Hall, David C. Martin, Vadim Nelidov, Theo Offerman, and Jeroen van de Ven. 2023. "Cognitive Biases: Mistakes or Missing Stakes?" *Review of Economics and Statistics*, 105(4): 818–832.

Hansen, Bruce E. 2024. "Jackknife Standard Errors for Clustered Regression." *Working Paper, the University of Wisconsin-Madison.*

Webb, Matthew D. 2023. "Reworking Wild Bootstrap Based Inference For Clustered Errors." *Canadian Journal of Economics*, 56(3): 839–858.