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

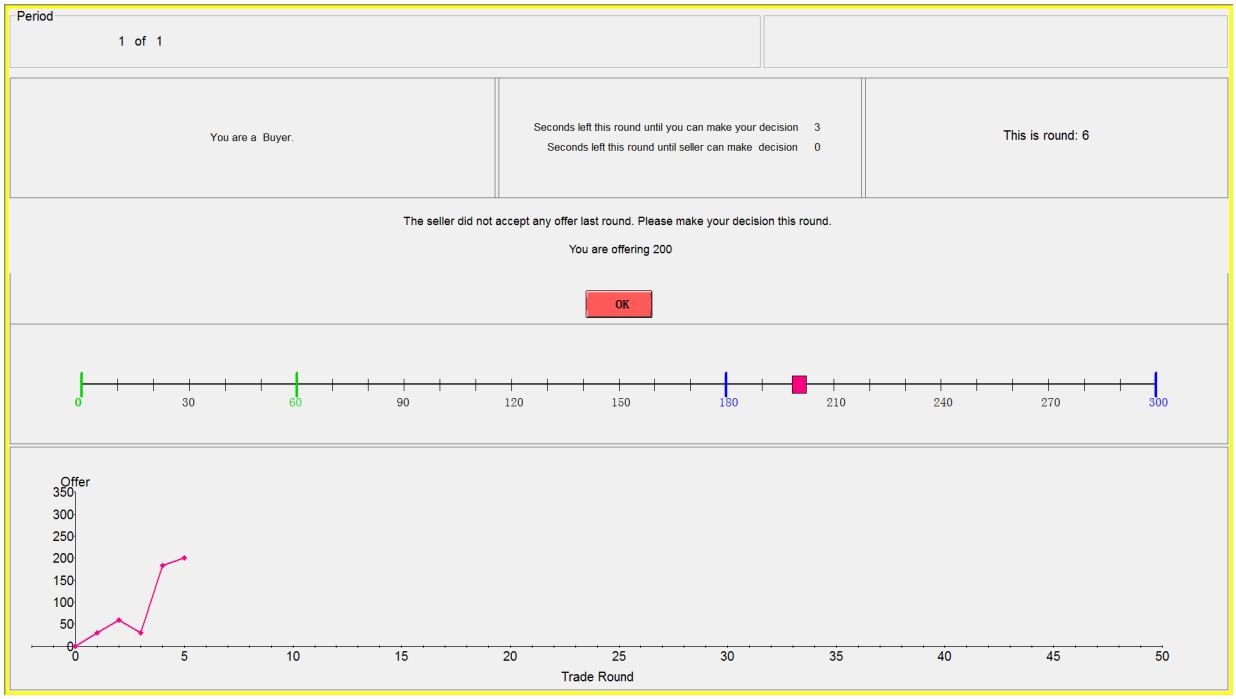

Figure A.2: Screenshot of Seller in the C1 Treatment

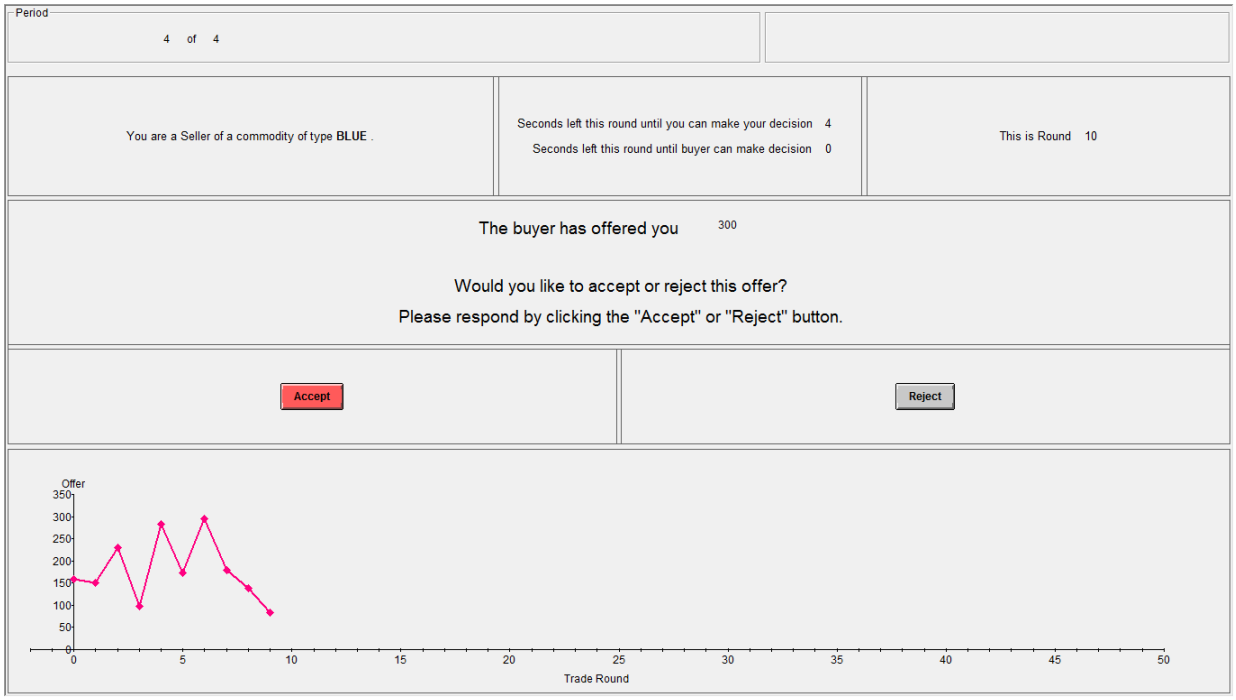

A.2 Instructions for Treatment B1L

. . . . . .

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

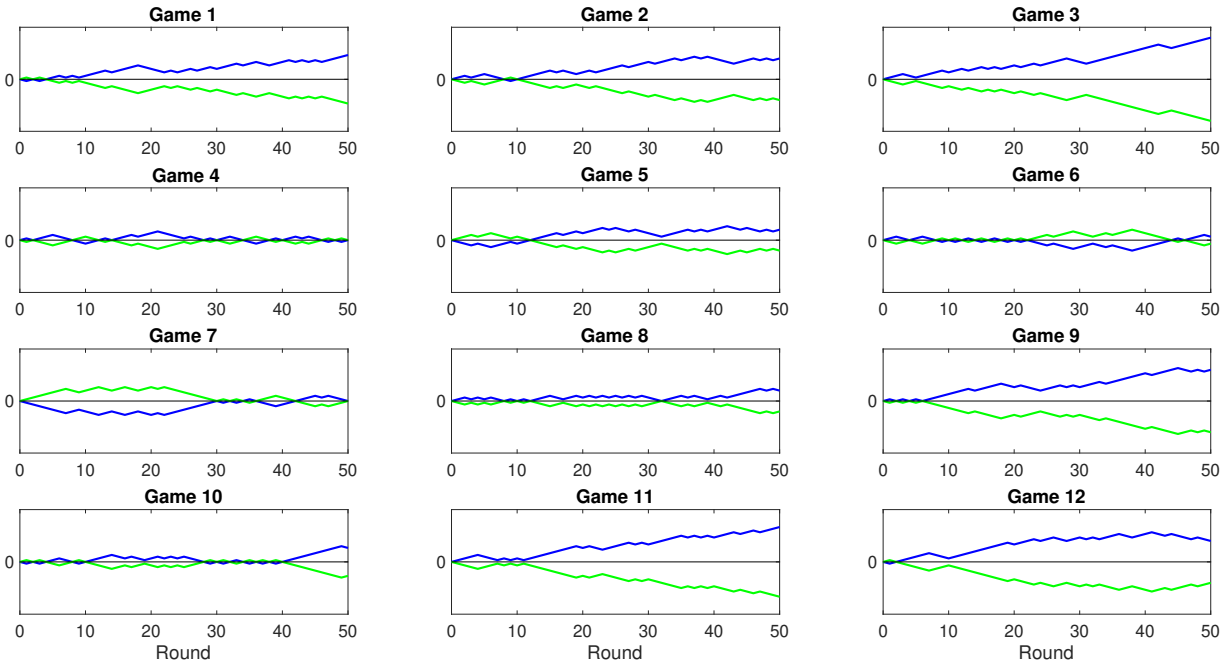

Figure A.4: News Displayed in the B1L Treatment

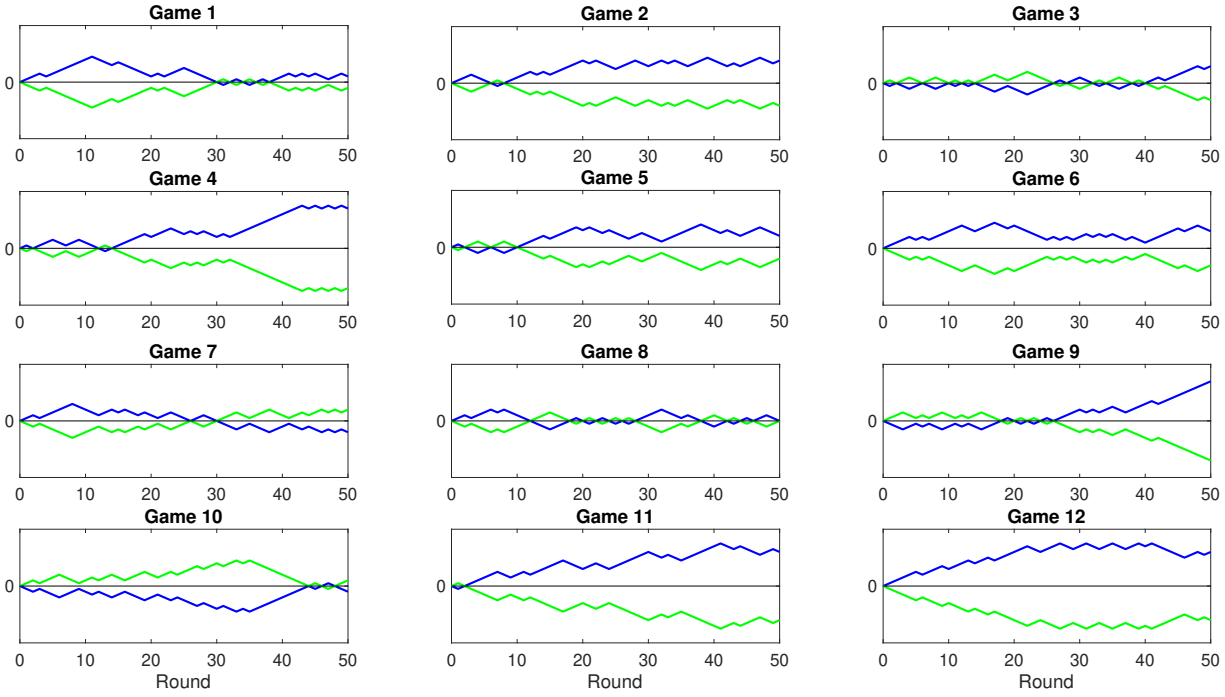

B Additional Empirical Results

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.