Rational Inattention in Games: Experimental Evidence

Submission for Experimental Economics' Special Issue in memory of Amnon Rapoport

David Almog* and Daniel Martin†

November 18, 2023

Abstract

To investigate whether attention responds rationally to strategic incentives, we experimentally implement a buyer-seller game in which a fully informed seller makes a take-it-or-leave-it offer to a buyer who faces cognitive costs to process information about the offer's value. We isolate the impact of seller strategies on buyer attention by exogenously varying the seller's outside option, which leads sellers to price high more often. We find that buyers respond by making fewer mistakes conditional on value, which suggests that buyers exert higher attentional effort in response to the increased strategic incentives for paying attention. We show that a standard model of rational inattention based on Shannon mutual information cannot fully explain this change in buyer behavior. However, we identify another class of rational inattention models that is consistent with this behavioral pattern.

*Kellogg School of Management, Northwestern University, david.almog@kellogg.northwestern.edu

†University of California, Santa Barbara and Kellogg School of Management, Northwestern University, danielmartin@ucsb.edu

1 Introduction

We are witnessing an era of exponential growth in the information accessible to consumers, and this overabundance of information is testing the limits of consumers' processing capabilities. For instance, to determine the quality of a product while shopping online, a consumer might need to evaluate a multitude of reviews, ratings, and technical specifications, and as a result, they may choose to make a purchase without being fully informed about the quality of all available products. Another common example is financial trading, where the consumers of sophisticated financial products face cognitive costs to process information about the value of these assets (e.g., the prospectuses of real estate collateralized debt obligations).

In these examples, consumers must decide how much attention to pay to the information that is available about product quality. But do consumers adjust their attention to account for the pricing strategies used by firms? If instead, consumer attention is invariant to strategic incentives, then this represents an opportunity for firms to take full advantage of the asymmetric and incomplete information produced by consumer inattention.[1]

We shed light on this question by studying whether individuals adjust their attention in response to the strategies of others, and if so, whether the adjustment they make is optimal. In other words, whether they are rationally inattentive in games. This question is very challenging to answer with market data because it is likely that many factors simultaneously change as firm strategies change. Thus, we attempt to answer this question by experimentally implementing a "buyer-seller" game in which the seller makes a take-it-or-leave-it price offer to the buyer after being fully informed about product value. The buyer sees this price clearly, but must perform a cognitive task to learn the product's value. We isolate the impact of seller strategies on buyer attention by varying only the seller's outside option and leaving the buyer's problem otherwise unchanged. In this game, the seller's outside option is the payoff the seller obtains when their offer is rejected by the buyer, which we exogenously vary between subjects. In real markets, variation in the outside option could be due the type of product (e.g., durable or not), the time of the year (e.g., holiday or not), or the market structure (e.g., many small firms or few large firms).[2]

We find that players in both roles meaningfully change their behavior when the seller's outside option changes, despite the subtlety of this between-subject treatment variation.

[1]Questions of this nature call for the integration of human cognition considerations into game theory. Professor Amnon Rapoport was an early advocate for incorporating the concepts of bounded rationality and information processing limitations into game theory (e.g., Rapoport 1999).

[2]The idea of outside options as a source of bargaining power is not new. Binmore, Shaked, and Sutton (1989) provide evidence on how a bargainer can use her outside option to gain leverage, and Weg, Zwick, and Rapoport (1999) build on this by showing how manipulation of the environment through outside opportunities can induce subjects to approximate rational behavior.

First, we find that sellers in the high outside option treatment (henceforth HOO) choose to set a high price more frequently. Second, buyer beliefs about seller pricing strategies shift across treatments in a similar direction and magnitude as the actual shift in seller pricing strategies. Third, buyers in the HOO treatment make fewer mistakes conditional on product value: they are less likely to reject a high-value product offer that is priced high or accept a low-value product that is priced high.

Does the change in buyer mistakes that we observe mean that subjects have adjusted their attention? Subjects can be modeled as making an attentional adjustment under the standard Bayesian expected utility framework if and only if buyer choices satisfy the *No Improving Action Switches (NIAS)* condition of Caplin and Martin (2015). If we take game payoffs at their face value, the NIAS condition is not satisfied in our experiment because subjects reject offers at a very high rate, so much so that there would be a expected utility gain from a wholesale switch from rejecting offers to accepting offers. However, this finding is consistent with a behavioral literature (e.g., Mitzkewitz and Nagel 1993; Camerer and Thaler 1995; Rapoport, Sundali, and Seale 1996) that demonstrates individuals reject apparently unfair offers in the Ultimatum Game when there is outcome uncertainty. We move forward by amending our theoretical framework to include a parameter that accounts for the disutility of accepting unfair offers, and we utilize the NIAS conditions in a novel way: to partially identify this behavioral parameter. Once we account for the disutility of accepting unfair offers, it becomes possible to model buyers as if they are adjusting their attention to the seller's outside option.

But is the buyer attentional adjustment optimal given their cognitive costs of attention? Answering this question requires a model of cognitive costs, and we begin by considering a leading model of rational inattention (RI) with costs that scale linearly with Shannon mutual information (Matějka 2015).[3] We find that the Shannon model is able to predict some of our empirical results, but it is not able to explain why individuals in the HOO treatment have a lower mistake rate in both states. This offers a larger insight about RI in games. As shown by Martin (2017), the Shannon model can be theoretically useful for analyzing games because optimal posterior beliefs are largely invariant to the prior, but it is precisely this property that is violated in our experiment.

On the other hand, optimal posterior beliefs can vary with the prior in RI models that are based on costly experiments (e.g., Denti, Marinacci, and Rustichini 2022; Pomatto, Strack, and Tamuz 2023).[4] In this second class of models, it is instead the costs of information that

[3] We follow Dean and Neligh (2019) in referring to any model where individuals choose their information to maximize expected utility subject to additive costs of information as an "RI" model, and to the RI with costs that scale linearly with Shannon mutual information (a measure of the entropy reduction between priors and posteriors) as the "Shannon model."

[4] In the language of Denti, Marinacci, and Rustichini (2022), "statistical experiments" map states of

are invariant to the prior. This property is sensible in settings where opponent strategies can impact the prior but not the costs of information, such as in game we implement experimentally. Indeed, buyer choices in our experiment pass a version of the *No Improving Action Switches (NIAC)* condition of Caplin and Dean (2015) that was adapted by de Clippel and Rozen (2021) to characterize this second class of models. This means that buyer behavior in our experiment is consistent with an RI model based on costly experiments.

1.1 Related Literature

Large literatures in economics, psychology, and neuroscience have provided evidence that individual decision making is influenced by cognitive limitations. A first-order question for economists is how these factors impact behavior in games (see Rapoport 1999 and Camerer 2003). One such cognitive limitation is the difficulty that individuals face in processing readily available information about payoff-relevant states, and RI is a leading approach to modeling how individuals trade off the costs of processing this information with its benefits.[5] Recently, a number of papers have considered the game theoretic implications of this modeling approach (e.g., Gentzkow and Kamenica 2014; Matějka 2015; Yang 2015; Martin 2017; Bloedel and Segal 2018; Boyacı and Akçay 2018; Lipnowski, Mathevet, and Wei 2020; Ravid 2020; Yang 2020; Albrecht and Whitmeyer 2023; Matyskova and Montes 2023).

However, most experimental results on RI are focused on non-strategic settings. Using individual decision making tasks, Dean and Neligh (2019) demonstrate in a robust way that subjects adjust their attention in response to changes in rewards.[6] At the same time, they find that subject behavior in their experiments is largely inconsistent with the Shannon model. Interestingly, the failures of the Shannon model differ between our experiment and theirs. The failure of the Shannon model in their experiment is that subjects increase their accuracy too slowly in response to changing incentives relative to the predictions of the Shannon model. Despite this, subjects in their experiment largely satisfy the foundational Locally Invariant Posteriors (LIP) condition, which is violated in our experiment. One possible explanation for this difference is that the attentional response of subjects might differ between individual decision making tasks and strategic settings. For instance, the behavioral force that we set-identify using NIAS is particular to games, and not surprisingly,

nature in signal distributions. The Shannon model places costs on distributions of posteriors, whereas this second class of models places costs on distributions of signals independently of the prior.

[5] See Caplin (2016) and Maćkowiak, Matějka, and Wiederholt (2023) for reviews of the RI literature, and Heidhues and Kőszegi (2018) for a review of important IO applications.

[6] Dewan and Neligh (2020) and Caplin, Csaba, Leahy, and Nov (2020) also run experiments in which there is attentional responsiveness to variation in incentives. Avoyan and Romagnoli (2023) allow individuals to choose their incentive level for an attentional task, which provides a novel way to measure attentional responsiveness.

NIAS passes in their experiment without the need to accommodate additional behavioral forces.

In another non-strategic experiment, Ambuehl, Ockenfels, and Stewart (2022) investigate who chooses to participate when incentives and attentional costs change. They find that the "revealed posteriors" of subjects (the observed frequencies of the good state conditional on the subject accepting or rejecting the gamble) are not Blackwell-ordered across any of the treatment variations, another distinction to our results that may stem from the strategic nature of our setting.

A related branch of the experimental literature is a growing set of papers that consider inattention to game payoffs. Brocas, Carrillo, Wang, and Camerer (2014) make use of an innovative mousetracking technique to uncover the reasons behind Nash play deviations in strategic games with private information. The player's choices and lookup data link deviations from Nash equilibrium play to a failure to consider all the necessary payoffs, which supports the theory of imperfect attention. Avoyan and A. Schotter (2020) study allocation of attention across games, and find that subjects choose to allocate their attention to games that have certain properties, like higher maximum payoffs. In a subsequent paper, Avoyan, Ribeiro, A. Schotter, E. Schotter, Vaziri, and Zou (2023) document a mismatch between planned and actual attention, using eye-tracking data to elaborate further on what features of the game payoffs are driving this inconsistency. Martin and Muñoz-Rodriguez (2022) study inattention to payoffs in the BDM mechanism (which is equivalent to a one-person second price auction) and find evidence that subjects do not fully internalize the payoff structure. Our paper differs from this line of papers in that we consider inattention to a payoff-relevant state instead of studying inattention to the payoffs of the game. Both forms of attention lead to uncertainty about payoffs, but have different applicability in terms of external contexts.

The two closest papers to ours are concurrent works by de Clippel and Rozen (2021) and Spurlino (2022). de Clippel and Rozen (2021) experimentally implement a sender-receiver game in which senders can choose to obfuscate the state by adding cognitive costs to learn the state, and they find that senders choose to strategically obfuscate.[7] Using an adapted version of NIAC, they find evidence of optimal perception adjustment in receivers' aggregate stochastic choice data for obfuscated messages. While there are clear differences in the structure of the game studied in their paper and ours, both papers share a common goal in understanding whether strategic inference can impact attention in games. The main difference lies in how strategic inference varies across treatments the experiments. In de Clippel and Rozen (2021), the sender should aim to obfuscate always in the opposing-interest state regardless of the treatment, and the buyer's prior should change across treatments

[7] Related to this, Jin, Luca, and Martin (2021a) show senders elect to shroud information by sending complex messages when the state is not favorable for them.

because the communication channel becomes more noisy. In our experiment, the sender's decision on how to price low-value products varies with the treatment, so changes in prior rises endogenously due to changes in seller payoffs, which is central to our question of interest.

Spurlino (2022) also shares the goal of understanding whether strategic inference can impact attention in games. In the experiment he implements, two players must decide whether or not to spend cognitive effort to learn the state before deciding to accept or reject a "deal," and the deal only goes through if both players accept it. He finds evidence that lack of sophistication is the primary reason why players are not able to correctly anticipate opponents' attention to the state. Some distinctive features of the game he studies, which set it apart from ours, include the fact that both players must exert cognitive effort to learn the state, and more importantly, the fact that the environment is altered through the cognitive costs of the other player.

2 Experimental Design

2.1 The Game

We begin by describing the "buyer-seller" game that we experimentally implemented. The seller is randomly assigned a hypothetical product that has a value to the buyer of either 100 or 50 Experimental Points (EPs).[8] Each value has an equal probability of being realized, which is known to both the buyer and the seller. The seller is shown the value of the product, and then offers to the buyer a price for that product of either 50 or 25 EPs. The buyer is told the product's price, and if he or she chooses to, the buyer had a way to learn the product's value by exerting cognitive effort before deciding whether accept or reject the seller's offer. If the buyer accepts the seller's offer, then the buyer gets an amount equal to the value of the product minus its price, and the seller gets an amount equal to its price. If the buyer rejects the offer, the buyer receives a fixed amount of 12.5 EPs and the seller received $\tau \in \{0, 20\}$ EPs, depending on the treatment. With these payoffs, (1) accepting a low price offer is a dominant action for the buyer, (2) accepting a high price and high value offer is better than the outside option, and (3) accepting a high price and low value offer is worse than the outside option.

The only difference between treatments is the seller's payoff when their offer is rejected, leaving the incentives on the buyer side unchanged. In the LOO treatment, sellers get $\tau = 0$ EPs if their offer is rejected. In the HOO treatment, the payoff of a seller whose offer is rejected increases to $\tau = 20$ EPs.

[8] Subject were told that EPs represented the probability of winning a €20 prize in a randomly selected round.

This game was implemented in two stages: a seller stage and a buyer stage. Every subject started with the seller stage, where they were asked to decide what their pricing strategy would be for the remainder of the session. Products with a value of 100 (the high value) were automatically priced at 50 (the high price), so subjects only had to select a strategy for pricing products with a value of 50 (the low value). To accommodate mixing, subjects were asked to provide a probability (from 0% to 100%) that the product would be given a price of 50 when the product had a value of 50 (it would be given a price of 25 with the remaining probability). We referred to this probability as a subject's "seller strategy" throughout the session.[9] To facilitate this decision, subjects selected their seller strategy using a slider, and the slider was accompanied by a table that was updated in real-time to show the probability of choosing each price given the current slider location.[10] Because products with a value of 100 (the high value) are automatically priced at 50 (the high price), we refer to the probability that the product is priced at 50 (the high price) when the product had a value of 50 (the low value) as the "mimicking rate" as it is rate at which low-value sellers are mimicking the price of high-value sellers.

The buyer stage featured 16 identically-structured rounds in which subjects were asked to accept or reject seller offers. In each round, the computer first randomly and anonymously matched each subject with another subjects' seller strategy. For each pairing, the computer randomly selected the value of the product, and if it was the low value (50), the computer then randomly selected a price with the probability given by the matched seller strategy. The subject was then informed of the price and could learn the value of the product by performing a task of adding up 20 numbers.[11] The numbers were determined by randomly drawing twenty numbers between -50 and 50 such that they added up to the product's value, and the exact realization was not known by sellers in advance.[12] A round of the buyer stage ended if the buyer accepted or rejected the subsequent offer within 120 seconds. If no decision was made in 120 seconds, the offer was automatically rejected.

2.2 Lab Implementation and Extra Tasks

Our pre-registered experiment was run in-person at the WISO Experimental Lab of Hamburg University.[13] A total of 12 sessions were conducted, and each session was composed of between 10 and 30 subjects, for a total of 238 subjects.[14] The experiment was organized and recruited using the hroot software package (Bock, Baetge, and Nicklisch 2014).

[9] See Figure 8 in Appendix A for an example screenshot.

[10] The slider started at a random location to control issues related to anchoring or default effects.

[11] See Figure 9 in Appendix A for an example screenshot.

[12] Subjects were not allowed to use scratch paper or a calculator.

[13] Preregistration plan available at https://aspredicted.org/wi5g4.pdf.

[14] 120 and 118 subjects in the LOO and HOO treatments respectively.

At the beginning of each session, subjects received a copy of the instructions, which were also read aloud before the experiment began.[15] To determine whether subjects understood key features of the game after the instructions were read aloud, we asked two types of unincentivized comprehension questions. First, to determine whether the incentives of the game were fully understood, we displayed the payoff table and asked subjects what the payoffs would be for every one of the 12 possible scenarios (2 roles x 3 price-value combinations x 2 buyer actions).[16] We provided the payoff table during these questions because the payoff table was present in both stages of the game and we wanted to measure whether they were able to understand the incentives conditional on having the payoff table available. 73.5% of subjects answered all 12 questions correctly, and 88.6% at least 10. We also checked on their comprehension of mixed strategies. In this question we showed each subject a seller strategy (drawn uniform random) and asked the subject how many times out of 100 they would expect to see a high price (50) if the product was low (50) for that given strategy. 75.2% of subjects answered this comprehension question correctly, while 11.4% did not submit an answer. If a subject answered any of our comprehension questions incorrectly, they were informed of the correct answer, but were allowed to continue.

The comprehension questions were followed by a practice stage in which subjects were given the opportunity to practice both as a seller and buyer without any payoff consequences. In this stage, subjects first entered a practice seller strategy and then completed four buyer practice rounds. There were two main difference between the buyer practice rounds and the regular buyer rounds, which we made clear to subjects. First, in the buyer practice rounds subjects faced their own practice seller strategy instead of another subject's seller strategy. Second, we made the seller strategy visible during the buyer practice rounds.

The practice stage was followed first by the seller stage and then the buyer stage. After these stages had been completed, subjects undertook a series of extra tasks that were incentivized. The aim of our first extra task was to determine whether there was a treatment effect in terms of subject beliefs. In the middle of the buyer stage (between rounds 8 and 9) and then again right after the buyer stage (after round 16), we elicited each subject's beliefs about the average seller strategy in their session.[17]

A primary aim for the subsequent extra tasks, as with the unincentivized comprehension questions, was to measure balance in comprehension and cognitive skills across treatments and to control for these factors if not. First we performed a "memory test" in which we presented two empty payoff tables, one for each role, and asked subjects to fill in the 6

[15]The instructions for the LOO treatment, which are nearly identical to the instructions for the HOO treatment, can be found in Appendix A.

[16]See Figure 7 in Appendix A for an example screenshot.

[17]To increase familiarity, we made the interface for belief elicitation similar to the interface from the seller stage.

empty spaces for each table without access to the actual payoff table. One purpose of this memory test was to determine whether subjects that fully internalized the payoff scheme and whether that was balanced across treatments. In addition, we used these questions to investigate the possibility that subjects in the experiment suffered from imperfect attention to game payoffs (as in Brocas, Carrillo, Wang, and Camerer 2014). In particular, we looked for evidence that subjects were more attentive to payoffs for the buyer role since they played that role for a longer portion of the experiment. We have two pieces of evidence that suggest this was not the case. First, in our comprehension questions on payoffs administered before the game, subjects responded similarly to questions about the payoff outcomes for both roles. Secondly, in the post-game memory test, we did not observe a significant difference in how they recalled the payoffs for one role versus the other.

We followed the memory test with a "math test", which was composed of 5 rounds. This task had the same format as the addition task in the buyer stage, but without seller provided prices, which had the purpose of identifying math skill separately from strategic decision making. For this task, we used 20 numbers from -100 to 100 (instead from -50 to 50) and asked subjects to provide an exact numeric answer, which was not constrained to be 50 or 100. Following the math test, we asked subjects how many questions they believed they answered correctly on the math test and the average number of questions they believed the other subjects in their session had answered correctly.

Every subject received an €8 show-up fee. In addition, they had the chance to earn two independent €20 rewards associated with their performance in each role, and up to €9.2 more through incentivized extra tasks. To translate performance into monetary prizes, a random round was selected for each of the roles and subjects received the EPs they earned in that round for the specific role. At the end of the session, two independent lotteries were realized, where the EPs earned as a seller represented the probability they received the first €20 reward, while the EPs earned as a buyer represented the probability of receiving the second €20 reward. The average payoff was €23.9 for an experiment that lasted approximately 75 minutes.[18] The experiment was programmed in z-Tree (Fischbacher 2007).

2.3 Design Choices

In this section, we discuss several of the main choices we faced while designing the experiment.

Two stages versus repeated interactions. We decided to split the experiment into two stages, beginning with a contingent strategy for sellers, so that we could have more observations of buyer behavior, which we felt was more likely to be stochastic given the attentional

[18]Payoffs averaged €21.2 in the LOO treatment and €26.5 in the HOO treatment, and ranged from €8 to €56.7.

task that they faced. In addition, understanding buyer behavior was the focus of our study, so we wanted subjects to be focused more on that role.

Human sellers versus computer sellers. Even though the focus of our study is the buyers side of this game, we used human sellers for two reasons. First, we felt that thinking through the seller's problem would help buyers better understand the game. Second, perceptions of computerized or algorithmic players is potentially in flux give the rapidly changing abilities (and subsequent news coverage) of advanced computer systems and artificial intelligence.

Payments. We pay subjects using "probability points" because it theoretically removes (under the assumption of expected utility theory) concerns about risk aversion. We decided to pay a random round for both of the two roles to incentivize effort throughout the whole experiment.

Number of rounds. In our pilot sessions, we asked subjects the round in which they started to experience fatigue, and we incorporated that feedback by reducing the practice rounds from 5 to 4 and the game rounds from 20 to 16. In addition, we introduced mid-game belief elicitation partly with the intention of giving subjects a break from repeatedly performing the same task. We also provided considerably ample time to sum the numbers in the buyer stage (120 seconds) in case fatigue slowed down some subjects. In fact, we found that performance actually improved, which suggests that fatigue effects were overshadowed by experience.

Eliciting the pricing strategy. We felt that narrowing the seller's pricing strategy to only one state would substantially reduce the complexity of the seller's problem. Further, in an initial version of this experiment (Martin 2016), sellers set a high price in 96.6% of high value. Thus, we felt that automatically setting a high price for these products would not result in a significant differences in seller behavior.

Practicing against themselves versus others. We felt that having subjects practice against themselves (their own seller strategy) would help them to more quickly understand the game and also to help them revisit their strategies in both roles. 86.6% of subjects reported in the follow-up questionnaire that the practice rounds helped them understand the game better. In addition, we did not want the practice rounds to be a channel for learning about the strategies of others.

Game parameter selection. We selected the game parameters with two objectives in mind: producing reasonable final payments and having good theoretical properties. While not directly comparable due to differences in payments (money versus probability points), our parameters produce roughly the same expected utility under the assumption of linear utility for money as the initial version of this experiment (Martin 2016) for the buyer and seller behavior in that experiment. The high prices were selected to produce an even split of the surplus to minimize fairness concerns.

Treatment parameter selection. When selecting the possible values for the seller's outside option, we first reduced the candidates to $\tau \in [0, 25)$. Negative values would have been incompatible with our payment mechanism, and values above 25 would have made it a dominant strategy for a seller to always set a high price. We selected 0 and 20 to produce as large a treatment effect as possible using round numbers.

Fixed time limit per round. We used a fixed time limit in each round for two reasons. First, to limit the overall duration of the experiment. Second, so that we could require all subjects to finish a round before any subject moved to the next one. The purpose of this design feature was to minimize the opportunity cost of time and because payments for the seller role required all other buyers to complete the experiment.

3 Treatment Effects

Table 1 provides a summary of subject characteristics across treatments and indicates whether the treatments are balanced with respect to these characteristics. This includes variables that reflect demographics, attention to the experimental instructions (comprehension questions and memory task), and math skill, which is our best proxy for cognitive skills. The only unbalanced attributes are the dummy variable for having a friend in the session and the math test score (especially in the extensive margin).[19]

3.1 Seller Strategies

Sellers chose to price high (price=50) when they had low-value products (value=50) at a higher rate on average in the HOO treatment than in the LOO treatment. In other words, sellers had a higher average seller strategy (mimicking rate) in the HOO treatment. The average seller strategies (mimicking rates) for the LOO and HOO treatments were 43.6% and 55.2% respectively, with a two-sided t-test p-value of 0.0054.[20] Thus, increasing the outside option from 0 to 20 EPs resulted in a 27% increase in the average mimicking rate. Figure 2 shows that a big part of the treatment effect on seller strategies is due to a shift in extreme strategies (defined as never or always mimicking). In the LOO treatment, 65.5% of sellers that used an extreme strategy never mimicked, while in HOO this number was 44.8%.

As a robustness check, we elicited within-subject measure of the treatment effect on seller strategies. After all extra tasks had been completed but right before the results of the

[19] Including a control for whether the subject has a friend in the room is also used in Jin, Luca, and Martin (2021b) and Jin, Luca, and Martin (2022).

[20] These are the chosen mimicking rates and not the actual mimicking rates realized in the experiment, which are of a similar magnitude.

	LOO		HOO		Comparison	
	Mean (sd.)	Obs.	Mean (sd.)	Obs.	Diff	t-stat
Male (dummy)	.401 (.49)	112	.5 (.50)	112	-.09	1.4
Age (years)	26.4 (5.9)	120	25.9 (5.8)	118	.55	0.7
English level (1-5)	4.04 (.84)	117	4.05 (.91)	118	-.01	0.07
Friend in session (dummy)	.116 (.35)	120	.042 (.20)	118	.074	2.1**
$\geq$ 10/12 on payoff comp. (dummy)	.891 (.31)	120	.881 (.32)	118	.010	0.24
12/12 on payoff comp. (dummy)	.733 (.44)	120	.733 (.44)	118	0	0.06
Mixed strategy comp. (dummy)	.741 (.40)	120	.762 (.42)	118	-.02	0.37
8/8 on memory task (dummy)	.633 (.48)	120	.703 (.45)	118	-.07	1.4
Math test correct (0-5)	.416 (.90)	120	.677 (1.17)	118	-.26	1.92*
Math test close (0-5)	.9 (1.12)	120	1.34 (1.40)	118	-.44	2.7***
$\geq$ 1/5 math test (dummy)	.258 (.43)	120	.338 (.47)	118	-.080	1.35
$\geq$ 1/5 math test close (dummy)	.533 (.50)	120	.618 (.48)	118	-.08	1.3

Notes: Observation is per subject. Value is missing if demographic information not provided by the subject.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: Subject characteristics and balance across treatments.

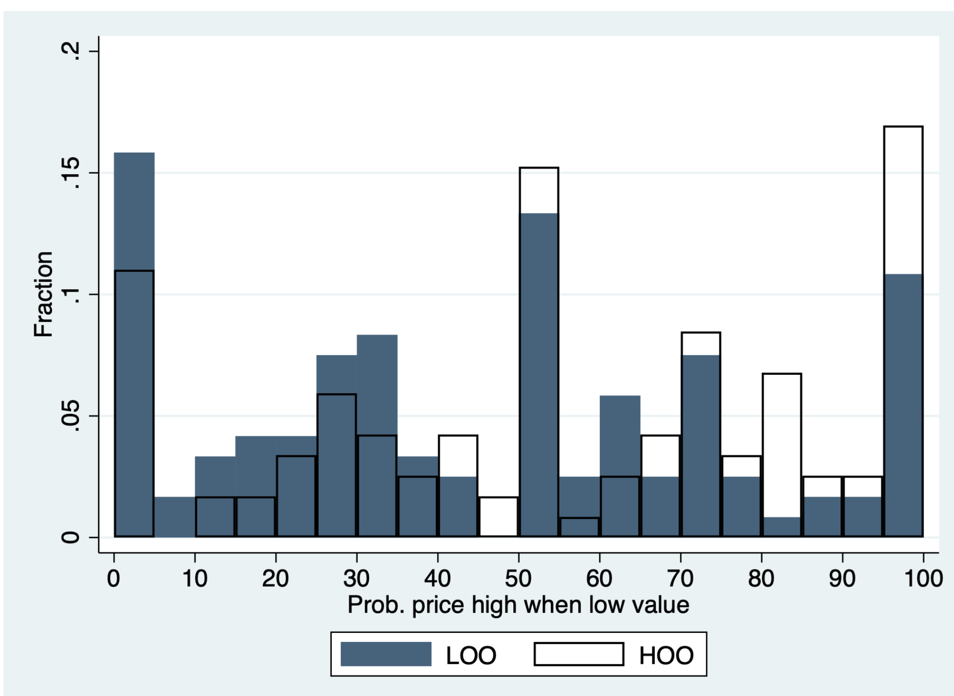

Figure 1: Distribution of seller strategies (mimicking rates) by treatment.

game were revealed, we asked subjects the following unincentivized question: "If you were to participate in exactly the same game, but the only difference was that the payoff sellers get when their offer is rejected is now [value in the other treatment] instead of [own treatment value], how would you change your seller strategy?" The possible answers were to increase, decrease, keep, or not sure.

We find that the between-subject treatment effect on seller strategies from the experiment is qualitatively validated by this self-reported within-subject measure. Figure 11 in Appendix B shows that a vast majority of the subjects in the LOO treatment reported that they would increase their seller strategy (price high more often) if they were playing in the alternative treatment, whereas most HOO subjects indicated they would like to decrease it. The fact that subjects in different treatments reacted in opposite directions relaxes potential concerns related to the framing of the question or a lack of incentives. The share of subjects that answered "same" and "not sure" is balanced across treatments.

3.2 Buyer Beliefs

Average elicited beliefs from the end of the buyer stage suggest that buyers in the HOO treatment believed there was more mimicking on average (61.8%) compared to buyers in the LOO treatment (50.8%). This difference is statistically significant at a 1% level for a two-sided t-test of equality of means (p= 0.001). Importantly, the treatment effect on beliefs goes in the same direction and is of a similar magnitude as the treatment effect on seller strategies.

One concern could be that beliefs are just mirroring each subject's own seller strategy. Figure 2 shows that even when we condition seller strategies, there is still an effect on beliefs for most bins (binned by groups of 10 percentage points). In an OLS regression that controls for own seller strategy, there is a sizeable and statistically significant relationship at the 5% level between elicited beliefs and treatment (the coefficient on HOO is 8.4 and has a p-value of 0.011).

Given that these beliefs were elicited at the end of the experiment is possible that some share of this effect is due to an experience factor for observing different frequencies of high prices across treatments. In line with this, the relationship between elicited beliefs in the middle of the buyer stage and treatment is smaller and no longer significant at a 5% level when controlling for own seller strategy (the coefficient on HOO is 4.2 and has a p-value of 0.0135). Given the dynamic change in beliefs, we will focus our analysis of buyer behavior on the second half of buyer stage rounds.

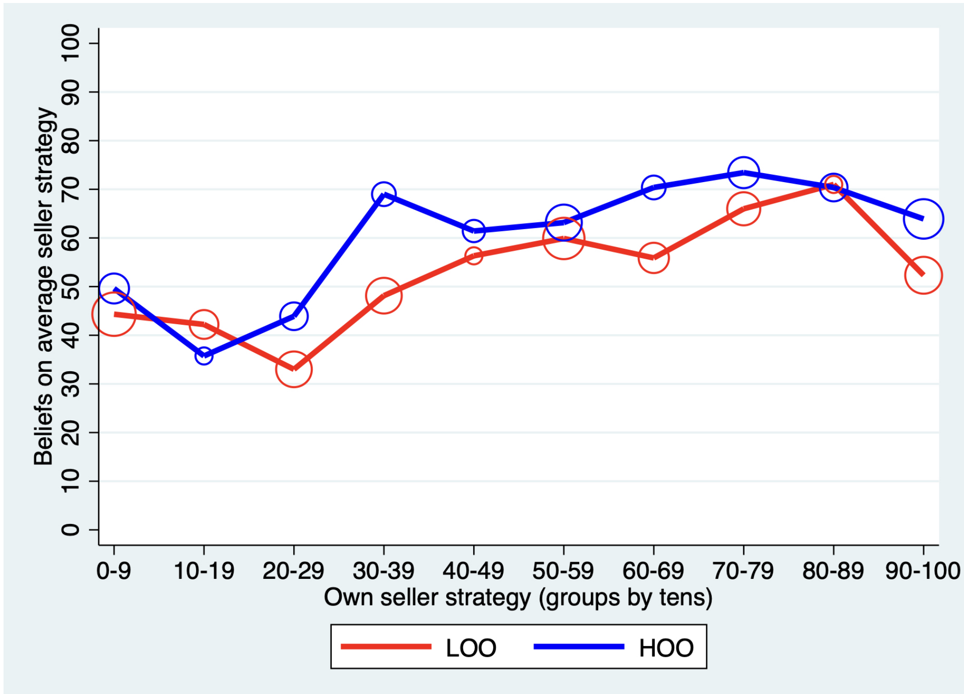

Figure 2: Treatment effect on elicited beliefs contingent on own seller strategy. Bubbles provide the frequency of each seller strategy bin. Elicited beliefs are from the end of the buyer stage.

3.3 Buyer Mistakes

In our experiment, buyers can make three types of "mistakes" (choices that are judged to be suboptimal ex-post). The first mistake, rejecting an offer with a low price, is the least common type of mistake, which is unsurprising due to the fact that it is a dominated action. Buyers observing a low-price offer should always accept the offer without the need to do the cognitive task, as the outside option gives a payoff of 12.5 and accepting a low-price offer gives a payoff of 25 or 75 depending on the value of the product. We use the choice of this dominated action as a way to test whether treatment and control are comparable in terms strategic reasoning. In the second half of rounds, the rates of rejecting low-price offers are 9.2% and 8.5% for the LOO and HOO treatments respectively, with a two-sided test of proportions p-value of 0.79.

The other two types of mistakes occur when buyers are facing high-price offers. These two mistakes, rejecting a high-value product and accepting a low-value one, will be the main focus of the analysis, as they will serve as the best proxy available to measure attention. Table 2 shows all three mistake rates for each treatment conditional on value and price. It reveals that subjects get better at avoiding the two undominated, high-price mistakes (the second and third columns of each block) in the second half of the experiment. While

	Rounds 1-8			Rounds 9-16			Overall		
(Value,Price)	50,25	50,50	100,50	50,25	50,50	100,50	50,25	50,50	100,50
LOO	5.8%	24.9%	23.6%	9.2%	18.2%	23.4%	7.5%	21.7%	23.5%
HOO	6.5%	23.9%	18.2%	8.5%	13.7%	17.8%	7.5%	18.5%	18.0%

Table 2: Mistake rates at each price and value by treatment.

this improvement does not rule out fatigue entirely, but it suggests that fatigue might be overshadowed by experience. To reduce dynamic effects in beliefs and mistakes, for all subsequent analyses we used only the second-half of rounds from the buyer stage (rounds 9 to 16) unless otherwise stated. That said, our main findings are robust to including all rounds.

Table 2 also shows that buyers facing high outside option sellers make fewer mistakes conditional on the state. Across treatments, there were 956 rounds where a high price was offered for high-value product, and subjects in the HOO treatment had a lower mistake rate of rejecting such offers: 23.4% versus 17.8% (with a two-sided test of proportions p-value of 0.033). There were 436 rounds where a high price was offered for a low-valued product, and subjects from the HOO treatment had a lower mistake rate of accepting such offers: 18.2% versus 13.7% (with a two-sided test of proportions p-value of 0.19) In this second case, the result is not significant, which could be due to a lack of statistical power.

Due to imbalance in the sample in the proportion of subjects having friends in the same session and in terms of performance on the math test, we use regression analysis to control for those factors. In Table 3 we pool together both types of high-price mistakes and regress them on a dummy variable for treatment and dummy variable controls for number of questions correct on the math test and for whether the subject had a friend in the session. In the regression specification with all controls included (the third and fourth columns), we find that subjects in the HOO treatment had a mistake rate of 4.74 percentage points lower (where the mean mistake rate was 21.4%). This result is significant at a 5% level without clustering, and when cluster standard errors at a subject level it is significant at a 10% level. The two controls we add seem quite relevant, as those that perform better on the math test and those who had a friend in the experiment had a significantly smaller mistake rates. The math performance is a conventional measure and was to be expected given the nature of the cognitive task in our experiment. However, finding that having a friend in the experiment has a statistically significant relationship suggests that it might be important to collect this variable in experimental games moving forward.

	(1)	(2)	(3)	(4)
	mistake	mistake	mistake	mistake
Seller's High Outside Option	-0.0528**	-0.0420**	-0.0474**	-0.0474*
	(0.0212)	(0.0213)	(0.0214)	(0.0285)
mathexact=1		-0.141***	-0.134***	-0.134***
		(0.0291)	(0.0293)	(0.0373)
mathexact=2		-0.144***	-0.149***	-0.149***
		(0.0378)	(0.0378)	(0.0340)
mathexact=3		-0.0987	-0.102	-0.102
		(0.0804)	(0.0804)	(0.0724)
mathexact=4		-0.0952	-0.0994	-0.0994
		(0.0732)	(0.0731)	(0.0721)
mathexact=5		-0.119	-0.125	-0.125
		(0.0782)	(0.0782)	(0.101)
Friend			-0.0759*	-0.0759*
			(0.0398)	(0.0399)
Cluster	None	None	None	Subject
N	1415	1415	1415	1415

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 3: Regressions of mistakes with high-price offers on treatment and imbalanced characteristics.

4 Attentional Adjustment?

Does this change in mistakes across treatments imply that buyers adjust their attention to available information in response to seller strategies? To investigate this possibility, we look to see if the choices that we have observed are consistent with a general model of inattention in which individuals adjust their attention to available information in response to other's strategies.

4.1 General Model of Inattention

Because there is a dominant action when prices are low, we focus on the case where prices are high. Conditional on seeing a high price, the prior probability of state $\theta \in \Theta$ is given by $\mu(\theta)$, and this depends on mimicking rate of sellers. In our experiment, the states are the product values $\Theta = [50, 100]$, and the probability of the high state, given by $\mu(100)$, decreases with the outside option because the average rate of mimicking increases with the outside option. It is this variation in the prior that drives the buyer's attentional adjustment in our model.

As in a large literature in information economics, we assume that the buyer chooses a Bayes-consistent information structure π, which stochastically generates posterior beliefs. For posterior $\gamma \in \Gamma$, $\gamma(\theta)$ is the probability of state θ in Θ. The buyer decides whether to buy the product or not depending on their posterior belief γ, and we assume the buyer maximizes expected utility when making a choice.

Technically, π is a function that maps the state into $\Delta(\Gamma)$, the set of probability distributions over Γ that have finite support, so that

$$\pi : \Theta \to \Delta(\Gamma)$$

Let Π denote the set of all such functions, $\pi(\gamma)$ be the unconditional probability of posterior $\gamma \in \Gamma$, $\pi(\gamma|\theta)$ be the probability of posterior γ given state θ, and $\Gamma(\pi) \subset \Gamma$ denote the support of a given π. We limit the set of information structures to those in $\Pi(\mu) \subset \Pi$ that generate correct posteriors for a given prior belief μ, so that

$$\Pi(\mu) = \left\{ \pi \in \Pi | \forall \gamma \in \Gamma(\pi), \forall \theta \in \Theta, \gamma(\theta) = \frac{\mu(\theta)\pi(\gamma|\theta)}{\sum\limits_{\theta \in \Theta} \mu(\theta)\pi(\gamma|\theta)} \right\}$$

The results of Caplin and Martin (2015) can be applied to show that buyer choices are consistent with this model of attentional adjustment if and only if they satisfy the *No Improving Action Switches (NIAS)* condition for both treatments. The NIAS condition requires there should be no utility gain from wholesale action switches in either treatment.

That is, neither a wholesale switch to accepting every time that buyers rejected, nor a wholesale switch to rejecting every time that buyers accepted should improve expected utility in either treatment. Normalizing the utility of the prize to 100 (so that one probability point is worth one expected utility unit), and denoting $P(\theta|reject)$ as the probability of state θ when buyers reject and $P(\theta|accept)$ as the probability of state θ when buyers accept,[21] this condition can be written as:

$$P(\theta_H|reject) * 12.5 + P(\theta_L|reject) * 12.5 \geq P(\theta_H|reject) * 50 + P(\theta_L|reject) * 0 \quad (1)$$

$$P(\theta_H|accept) * 50 + P(\theta_L|accept) * 0 \geq P(\theta_H|accept) * 12.5 + P(\theta_L|accept) * 12.5 \quad (2)$$

Despite the apparent weakness of this test, we find that this condition is not satisfied on aggregate for either treatment. Thus, it does not appear that buyer choices in our experiment are consistent with a model in which attention is adjusting in response to sender strategies. However, this analysis ignores a well-established behavioral effect in the literature, which we add to our model in the following section.

4.2 Disutility of Accepting Unfair Offers

In the literature on Ultimatum Games, there is a version of this classic game, called the Demand (Ultimatum) Game, in which the pie is of an uncertain size. For this game, it has been shown that individuals will reject seemly unfair offers, even if means receiving no return (see Mitzkewitz and Nagel 1993, Camerer and Thaler 1995, and Rapoport, Sundali, and Seale 1996).[22] To identify the disutility of accepting unfair offers in terms of probability points, we again draw on the NIAS condition. As mentioned before, this condition requires that there should be no utility gain from wholesale action switches (a wholesale switch from rejecting to accepting and vice versa).

Results on the Demand (Ultimatum) Game suggest that the expected utility from accepting an unfair offer might not be 0 (the number of probability points received), but instead, be well below this because of additional psychological costs. To recover the implied disutility of accepting an unfair offer, which we denote as ϕ, we instead ask what value of ϕ satisfies the NIAS condition:

$$P(\theta_H|reject) * 12.5 + P(\theta_L|reject) * 12.5 \geq P(\theta_H|reject) * 50 + P(\theta_L|reject) * \phi \quad (3)$$

$$P(\theta_H|accept) * 50 + P(\theta_L|accept) * \phi \geq P(\theta_H|accept) * 12.5 + P(\theta_L|accept) * 12.5 \quad (4)$$

[21] Caplin and Martin (2015) refer to this type of choice data as "state-dependent stochastic choice data".

[22] In other bargaining environments, such as the Tunisian Bazaar mechanism, Rapoport, Erev, and Zwick (1995) present comparable evidence, indicating that buyers appeared to "punish" sellers by rejecting offers made "too late". In this context, the sensation of unfairness arises from the slow decrease in prices.

This reduces to:

$$\frac{P(\theta_H|accept)*50-12.5}{P(\theta_H|accept)} \geq \phi \geq \frac{P(\theta_H|reject)*50-12.5}{P(\theta_H|reject)} \tag{5}$$

Looking first at state-dependent stochastic choice data for buyers in the LOO treatment, $P_{LOO}(\theta_H|accept)$ and $P_{LOO}(\theta_H|accept)$, the bounds on ϕ are -14.7 to -387.9. The state-dependent stochastic choice data for buyers in the HOO treatment, $P_{HOO}(\theta_H|accept)$ and $P_{HOO}(\theta_H|accept)$, produce wider bounds: -2.5 to -424.3.

To tighten the bound on this parameter we assume that subjects who performed well in the math test have the same ϕ as the subjects that did not perform well in the math test. The justification for this assumption is that since the math test is not a game, this behavioral factor, which is related to strategic play, might orthogonal to performance on the test. Looking just at those who performed poorly at the math test (had no problems correct on the math test), the estimated ϕ is tighter than before: -21.25 to -289.91937.

If we incorporate into our model a value for the disutility of accepting unfair offers in this range, the NIAS condition passes, which means that we *can* model individuals as if they are adjusting their attention in response to changes in opponent strategies. In other words, the mistake reduction that we found is evidence in favor of an attentional reaction from the buyers in response to varying the outside option of the sellers, as long as we account for the disutility of accepting unfair offers. What we want to do next is investigate whether the buyers are adjusting their attention optimally in response to changes in seller strategies.

5 Optimal Attention Adjustment?

To test whether buyers are adjusting their attention optimally in response to changes in seller strategies, we begin by considering the Shannon model, a leading model of RI with costs based on Shannon mutual information. We generate several theoretical predictions for how subject behavior would change in our game if buyers had attention costs of this form. For purposes of this section we use a value of $\phi = -50$ to reflect the disutility of accepting unfair offer, but all of the predictions are robust to using other values that satisfy the NIAS conditions.

5.1 Theoretical Framework

In the Shannon model a posterior is more costly if it reduces more uncertainty. Formally, each information structure $\pi \in \Pi(\mu)$ has a cost in expected utility units that is determined by the function

$$K(\pi,\lambda,\mu)=\lambda\left(\left[\sum_{\gamma\in\Gamma(\pi)}\pi(\gamma)\sum_{\theta\in\Theta}[\gamma(\theta)\ln(\gamma(\theta))]\right]-\sum_{\theta\in\Theta}[\mu(\theta)\ln(\mu(\theta))]\right)$$

where $\lambda\in\mathbb{R}_{++}$ is a linear cost parameter that is interpreted as the marginal cost of attention. In the binary state case, this functional form produces a u-shaped cost for each posterior, which increases symmetrically towards being certain of the state θ. In addition, an information structure that returns the prior μ as the posterior belief has a cost of zero. In other words, if there is no attention (and thus no information), then there is no attentional cost.

To make predictions in our setting for these assumptions about the cost of attention, we leverage the theoretical characterization of equilibrium from Martin (2017) for a buyer-seller game featuring RI with Shannon mutual information costs. We expand his model by incorporating parameters for the buyer and sellers outside options, τ and β respectively (buyer and seller payoffs in case the offer is rejected). In Theorem 1 of that paper, Martin (2017) establishes that in the most informative equilibrium, regardless of the marginal cost of attention λ, sellers of high-value products should price high and sellers of low-value products should mimic them with a positive probability. After incorporating parameters for the buyer and seller outside options, Theorem 1 in that paper can be restated as follows.

Theorem 1 *(Martin 2017) For any game G, there exists a mixed strategy PBE in which high value sellers set a high price p_H with probability 1 and low value sellers set price p_H with a unique probability η and otherwise set a low price p_L. Formally, this equilibrium is defined as:*

- *Seller pricing strategy $\hat{\sigma}$ (probability of setting each price for each value):*
  $\hat{\sigma}(p_H|\theta_H)=1,\ \hat{\sigma}(p_H|\theta_L)=\eta,\ \hat{\sigma}(p_L|\theta_L)=1-\eta$
- *Buyer attention strategy $\hat{\pi}$ (probability of each posterior for each price):*
  $\hat{\pi}_{p_H}(\gamma^1)=min\left\{\frac{\hat{\mu}(\theta_H|p_H)-\gamma^0(\theta_H)}{\gamma^1(\theta_H)-\gamma^0(\theta_H)},1\right\},\ \hat{\pi}_{p_H}(\gamma^0)=1-\hat{\pi}_{p_H}(\gamma^1),\ \hat{\pi}_{p_L}(\gamma^*)=1$
- *Buyer purchasing strategy $\hat{\alpha}$ (probability of accepting at each price and posterior):*
  $\hat{\alpha}(\gamma^1,p_H)=1,\ \hat{\alpha}(\gamma^0,p_H)=0,\ \hat{\alpha}(\gamma^*,p_L)=1$
- *Buyer prior beliefs $\hat{\mu}$ (beliefs about value for each price):*
  $\hat{\mu}(\theta_H|p_H)=\frac{\delta(\theta_H)}{\delta(\theta_H)+(1-\delta(\theta_H))\eta},\ \hat{\mu}(\theta_H|p_L)=0$
- *Optimal posteriors for low-price offers γ^* and when rejecting γ^0 and accepting γ^1 high-price offers:*

$$\gamma^{*}(\theta_H) = 0,\ \gamma^{0}(\theta_H) = min\left\{ \frac{exp\left(\frac{\beta}{\lambda}\right) - exp\left(\frac{\theta_L - p_H}{\lambda}\right)}{exp\left(\frac{\theta_H - p_H}{\lambda}\right) - exp\left(\frac{\theta_L - p_H}{\lambda}\right)}, \hat{\mu}\left(\theta_H | p_H\right) \right\},$$

$$\gamma^{1}(\theta_H) = max\left\{ \frac{exp\left(\frac{\theta_H - p_H}{\lambda}\right)}{exp\left(\frac{\beta}{\lambda}\right)} \gamma^{0}(\theta_H), \hat{\mu}\left(\theta_H | p_H\right) \right\}$$

- *Low-value seller mimicking rate* η*:*

$$\eta = min\left\{ \frac{\delta(\theta_H)}{1-\delta(\theta_H)} \frac{\left(1-\gamma^{0}(\theta_H)\right)\left(1-\gamma^{1}(\theta_H)\right)}{\gamma^{0}(\theta_H)(1-\gamma^{1}(\theta_H)) + \frac{p_L - \tau}{p_H - \tau}(\gamma^{1}(\theta_H) - \gamma^{0}(\theta_H))}, 1 \right\}$$

5.2 Prediction for Seller Strategies

This characterization result provides a map between interior mimicking rates and marginal costs of attention in equilibrium, which is illustrated in Figure 3 for the parameter values in the two treatments of our experiment. Figure 3 in Appendix B has an additional axis for treatment, which makes it more visually salient how the map gets steeper as the outside option increases.

This figure illustrates that low-value sellers in HOO treatment will choose to price high more often in equilibrium, regardless of the buyer's marginal cost of attention. Formally, introducing variation in the seller's outside option τ results in the following comparative static when $\eta < 1$:

$$\frac{\partial \eta}{\partial \tau} = \frac{\delta(\theta_H)}{1-\delta(\theta_H)} \frac{\left(1-\gamma^{0}(\theta_H)\right)\left(1-\gamma^{1}(\theta_H)\right)\frac{p_H - p_L}{(p_H - \tau)^2}\left(\gamma^{1}(\theta_H) - \gamma^{0}(\theta_H)\right)}{\left(\gamma^{0}(\theta_H)\left(1-\gamma^{1}(\theta_H)\right) + \frac{p_L - \tau}{p_H - \tau}\left(\gamma^{1}(\theta_H) - \gamma^{0}(\theta_H)\right)\right)^2} > 0 \tag{6}$$

Intuitively, the higher outside option leads low-value sellers to roll the dice more often, as being rejected is not as bad an outcome as compared to sellers in the LOO treatment. This generates our first testable theoretical prediction.

Prediction 1: Seller Behavior. Sellers in the HOO treatment will mimic at a *higher* rate.

In our experiment, seller behavior is consistent with this first prediction, as the average mimicking rates raise from 43.6% to 55.2% when we increase the outside option.

5.3 Predictions for Buyer Mistakes

Under the assumption that we are on an mixing equilibrium such that $\eta < 1$, increasing the seller's outside option τ results in the comparative static prediction that buyers will accept offers less often:

$$\frac{\partial \hat{\pi}_{p_H}\left(\gamma^{1}\right)}{\partial \tau} = \frac{1}{\gamma^{1}(\theta_H) - \gamma^{0}(\theta_H)} \frac{\partial \hat{\mu}\left(\theta_H | p_H\right)}{\partial \tau} < 0 \tag{7}$$

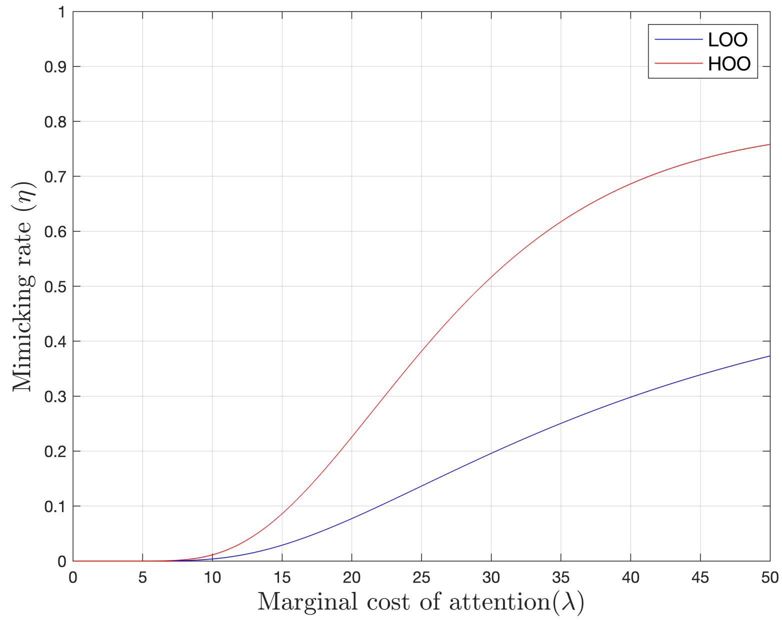

Figure 3: Theoretical relationship between mimicking and marginal cost of attention.

However, this does not tell us what the model predicts how buyer mistakes with change with the seller's outside option.

Panel a of Figure 4 shows that in equilibrium, the mistake of accepting a low-value product is lower in the HOO treatment whenever the mistake rate is positive. This is sensible given that buyers are accepting offers less often unconditionally. This same ordering is true if buyers best respond to the empirical distribution of seller strategies in our experiment, as show in Panel b of Figure 4. Because it is a partial equilibrium result, the latter ordering would also hold for any possible population distribution of cognitive costs λ. These two results lead us to our second main prediction.

Prediction 2: Mistake of accepting a low-value product when prices are high. Buyers in the HOO treatment will have a *lower* mistake rate of this type.

Intuitively, the second prediction points out that buyers in the HOO treatment should be better at rejecting low-value, high-price offers in the HOO treatment. The experimental evidence and theoretical predictions for the Shannon model are also aligned for this prediction, as buyers made this type of mistake for 18.2% of low-value, high-price offers in the LOO treatment to 13.7% of low-value, high-price offers in the HOO treatment.

Turning to the other kind of buyer mistake for high-price offers, panel a of Figure 5 shows that in equilibrium, the mistake of rejecting a high-value, high-price offer is higher in the HOO treatment whenever the mistake rate is positive. Once again, the same ordering is true if buyers best respond to the empirical distribution of seller strategies in our experiment, as show in Panel b of Figure 5. This ordering would also hold for any possible population

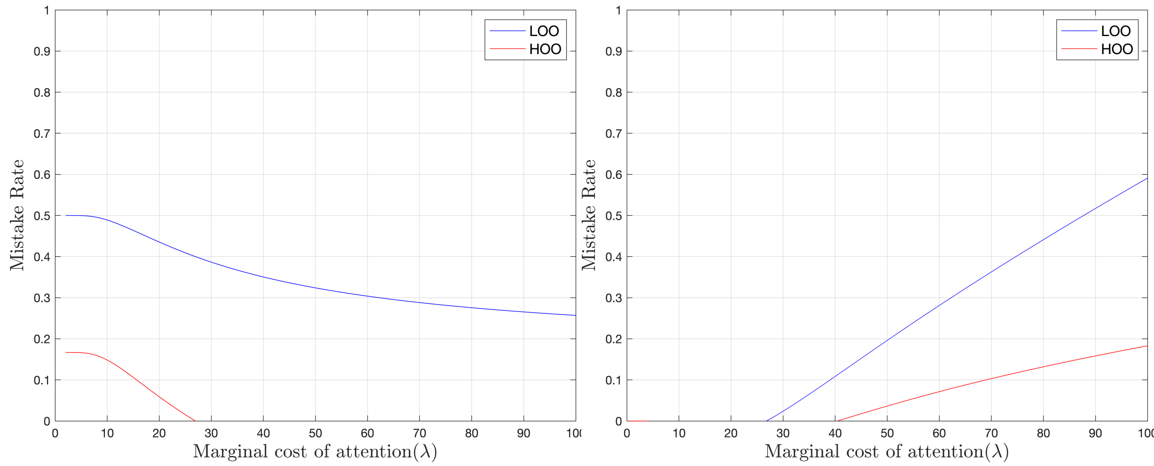

(a) Equilibrium predictions.

(b) Empirical best response.

Figure 4: Probability of accepting a low-value product at a high price.

distribution of cognitive costs λ. These two results generate our third main prediction.

Prediction 3: Mistake of rejecting a high-value product when prices are high. Buyers in the HOO treatment will have a *higher* mistake rate of this type.

After having the first two predictions from the model supported by the experimental data, it is the third prediction that moves in an opposite direction to what we observe in our experiment. We find that subjects in the HOO treatment rejected high-price, high-value offers less often in the HOO treatment (17.8%, compared to 23.4% in the LOO treatment).

5.4 Departures the Shannon Model

A natural question is why the subjects in our experiment are not behaving in line with the third prediction of the Shannon model. As established by Caplin and Dean (2013) and Caplin, Dean, and Leahy (2019), one of the consequences of using Shannon mutual information as the specific cost function is that the *Invariant Likelihood Ratio (ILR)* condition has to be satisfied for chosen actions. As we only have two actions in this context, the *ILR* condition boils down to:

$$\frac{\gamma^1(\theta)}{\gamma^0(\theta)} = e^{\frac{U(accept,\theta)-U(reject,\theta)}{\lambda}} \tag{8}$$

Given that the utility of accepting and rejecting offers does not change for subjects across treatments, nor should the marginal costs of attention λ, ILR implies that optimal posteriors should not change across treatments. This feature of the Shannon model has been shown to be very helpful in game-theoretic analysis (Martin 2017), but it also has very strong implications about how subjects should behave across treatments. Given that

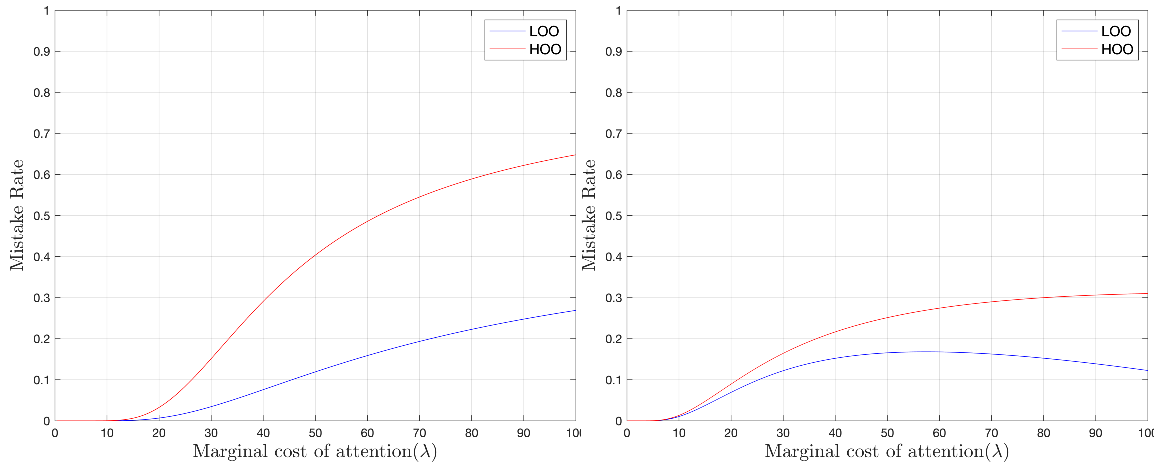

(a) Equilibrium predictions.

(b) Empirical best response.

Figure 5: Probability of rejecting high-value product at a high price.

optimal posteriors do not change with the sellers outside option, buyers should be equally well informed when accepting and rejecting for both outside option levels.[23] This is why the model predicts that increasing the outside option should lead buyers to reject high-price offers more often regardless of the product's value. But buyers in our experiment appear to be better informed when both accepting and rejecting. Figure 6 shows that the "revealed" posteriors (average likelihood at each action) are Blackwell-ordered across treatments (a welfare ordering studied by Caplin and Martin 2021).[24] Even though the prior probability of high quality is lower for high-priced offers in the HOO treatment, and subjects believe this to be the case,[25] this belief is reversed before accepting and their revealed posterior when accepting is higher (.921) than LOO subjects (.914). Likewise, when HOO subjects reject, their revealed posterior is also more accurate (.286) than the one from the LOO treatment (.421).

5.5 Prior Variation

Under the Shannon model, the buyer's optimal posteriors do not change with the seller's outside option, but the buyer's attention strategy does change because changes in the prior

[23]The only exception would be if subjects chose to be completely inattentive in one treatment because the seller was mixing at a very high or very low rate in that treatment.

[24]Building directly from the *ILR* condition, posteriors across treatments should overlap for both actions under the Shannon model.

[25]Using the elicited beliefs of the average seller strategy we can calculate the prior, which is how often subjects expect to receive a high-value offer conditional on observing a high price. The prior in the HOO treatment is lower (.636) compared to the LOO (.684), because buyers correctly anticipate more mimicking and observing a high price becomes a less informative signal of high value as mimicking increases.

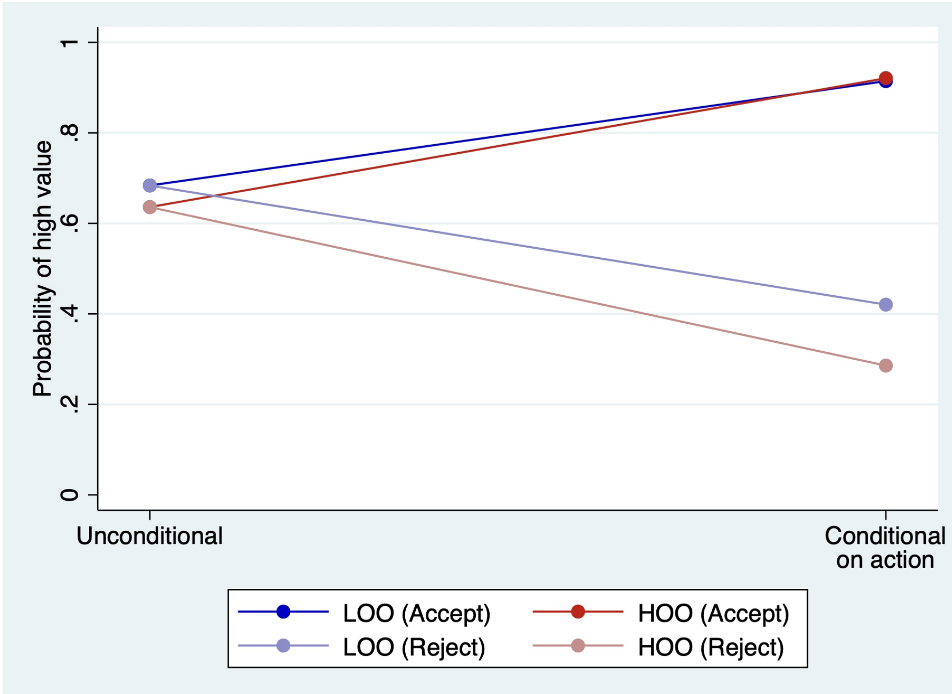

Figure 6: Probability of high value conditional on action.

(due to changes in seller pricing strategies) impacts the probability that buyers reach each posterior. However, the buyer's purchasing strategy does not change, it remains always accepting at one posterior and rejecting at the other.

An alternative model is that the attention strategy of buyers does not change with the prior, but that they change their purchasing strategy with the prior, such as changing their cutoff belief at which they take a particular action. However, the same pattern that invalidates Shannon model, the Blackwell ordering of revealed posteriors across treatment, also invalidates this alternative theory. Under this alternative model, one would expect an increase in rejections, even in cases where the state is high. This increase in rejections should correspond to more rejecting of high-prices, high-value offers, but that is not what we observe.

6 Attention as Costly Experiments

While optimal posterior beliefs are invariant to the prior with the Shannon model, optimal posteriors can vary with the prior in models of costly experiments (e.g., Denti, Marinacci, and Rustichini 2022; Pomatto, Strack, and Tamuz 2023). Unlike the Shannon model provided earlier, the costs in this class of models are not on the posteriors themselves. As noted in Denti, Marinacci, and Rustichini (2022) and Pomatto, Strack, and Tamuz (2023), this is advantageous to use in strategy settings in which one player can impact the prior over states, which is unlikely to impact the attentional costs of the opponent.

Buyer choices are consistent with this model if they pass an adapted version of the NIAC condition (Caplin and Dean 2015). de Clippel and Rozen (2021) note that this class of models can be characterized by the NIAC conditions if we treat the states as equiprobable, but adjust the utility to account for any prior variation.

To express this condition formally in our setting, we denote P_{LOO} and P_{HOO} as the state-dependent stochastic choice data of buyers in the LOO treatment and HOO treatment, respectively. This adapted version of NIAC requires that there is no way to cycle the state-conditional marginal distributions of P_{LOO} and P_{HOO} such that expected utility is increased:

$$\begin{aligned} &P_{LOO}(reject|\theta_H) * P_{LOO}(\theta_H) * 12.5 + P_{LOO}(reject|\theta_L) * P_{LOO}(\theta_L) * 12.5 \\ &+P_{LOO}(accept|\theta_H) * P_{LOO}(\theta_H) * 50 + P_{LOO}(accept|\theta_L) * P_{LOO}(\theta_L) * \phi \\ &-[P_{HOO}(reject|\theta_H) * P_{LOO}(\theta_H) * 12.5 + P_{HOO}(reject|\theta_L) * P_{LOO}(\theta_L) * 12.5 \\ &+P_{HOO}(accept|\theta_H) * P_{LOO}(\theta_H) * 50 + P_{HOO}(accept|\theta_L) * P_{LOO}(\theta_L) * \phi] \\ &+P_{HOO}(reject|\theta_H) * P_{HOO}(\theta_H) * 12.5 + P_{HOO}(reject|\theta_L) * P_{HOO}(\theta_L) * 12.5 \\ &+P_{HOO}(accept|\theta_H) * P_{HOO}(\theta_H) * 50 + P_{HOO}(accept|\theta_L) * P_{HOO}(\theta_L) * \phi \\ &-[P_{LOO}(reject|\theta_H) * P_{HOO}(\theta_H) * 12.5 + P_{LOO}(reject|\theta_L) * P_{HOO}(\theta_L) * 12.5 \\ &+P_{LOO}(accept|\theta_H) * P_{HOO}(\theta_H) * 50 + P_{LOO}(accept|\theta_L) * P_{HOO}(\theta_L) * \phi] \\ &\geq 0 \end{aligned}$$

The first two lines of this expression give the utility in the LOO treatment under the signal structure used that treatment, and the following two lines give the utility in the LOO treatment if instead we use the signal structure from the other treatment. The fourth and fifth lines give the utility in the HOO treatment under the signal structure used in that treatment, and the following two lines give the utility HOO treatment if we instead use the signal structure from in the other treatment.

For the P_{LOO} and P_{HOO} observed in our experiment, this condition reduces to:

$$-1.62852 + 0.0137712 * \phi + 1.549155 - 0.0161268 * \phi \geq 0 \tag{9}$$

This expression is satisfied; for example, if we use a ϕ of -50, as before.

7 Conclusion

In this paper we present the results of an experimental implementation of a "buyer-seller" game in which we varied the outside option of the sellers, keeping the buyer's problem otherwise intact. Increasing the outside option of a seller embodies an upsurge in their leverage against the buyer. This manipulation does not affect the buyer's payoff directly but it does encourage sellers to revise their strategy upward, so that they set high prices more often when selling low value products. Our central result shows that buyers foresee this mimicking upsurge and react by increasing their attention. Buyers make fewer mistakes regardless of the value, which to the best of our knowledge is the first evidence of an attentional response in a strategic environment triggered solely by inference about the opponent strategy. We find that the attentional reaction from the buyers is *not* consistent with the Shannon model (a model of RI with Shannon mutual information costs), which is a popular choice for modeling costly information acquisition due to its tractability. However, we find that buyer behavior *is* consistent with optimal attention models based on costly experiments.

References

Albrecht, Brian C and Mark Whitmeyer (2023). “Comparison shopping: Learning before buying from duopolists”. *arXiv preprint arXiv:2302.06580*.

Ambuehl, Sandro, Axel Ockenfels, and Colin Stewart (2022). “Who Opts In? Selection and Disappointment through Participation Payments”. *Review of Economics and Statistics*, pp. 1–45.

Avoyan, Ala and Giorgia Romagnoli (2023). “Paying for inattention”. *Economics Letters* 226, p. 111114.

Avoyan, Ala and Andrew Schotter (2020). “Attention in games: An experimental study”. *European Economic Review* 124.

Avoyan, Ala et al. (2023). “Planned vs. Actual Attention”. *Management Science*, pp. 1–22.

Binmore, Ken, Avner Shaked, and John Sutton (1989). “An Outside Option Experiment”. *The Quarterly Journal of Economics* 104.4, pp. 753–770.

Bloedel, Alexander W and Ilya R Segal (2018). “Persuasion with rational inattention”. *Available at SSRN 3164033*.

Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch (2014). “hroot: Hamburg Registration and Organization Online Tool”. *European Economic Review* 71, pp. 117–120.

Boyacı, Tamer and Yalçın Akçay (2018). “Pricing when customers have limited attention”. *Management Science* 64.7, pp. 2995–3014.

Brocas, Isabelle et al. (2014). “Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games”. *Review of Economic Studies* 81, pp. 944–970.

Camerer, Colin (2003). “Behavioral Game Theory: Experiments in Strategic Interaction”. Princeton: Princeton University Press.

Camerer, Colin and Richard Thaler (1995). “Anomalies: Ultimatums, Dictators and Manners”. *Journal of Economic Perspectives* 9.2, pp. 209–219.

Caplin, Andrew (2016). “Measuring and modeling attention”. *Annual Review of Economics* 8, pp. 379–403.

Caplin, Andrew and Mark Dean (2013). *Behavioral implications of rational inattention with shannon entropy*. Tech. rep. National Bureau of Economic Research.

— (2015). “Revealed Preference, Rational Inattention, and Costly Information Acquisition”. *American Economic Review* 105.7, pp. 2183–2203.

Caplin, Andrew, Mark Dean, and John Leahy (2019). “Rational inattention, optimal consideration sets, and stochastic choice”. *The Review of Economic Studies* 86.3, pp. 1061–1094.

Caplin, Andrew and Daniel Martin (2015). “A Testable Theory of Imperfect Perception”. *The Economic Journal* 125.February, pp. 184–202.

Caplin, Andrew and Daniel Martin (2021). “Comparison of Decisions under Unknown Experiments”. *Journal of Political Economy* 129.11, pp. 3185–3205.

Caplin, Andrew et al. (2020). “Rational inattention, competitive supply, and psychometrics”. *The Quarterly Journal of Economics* 135.3, pp. 1681–1724.

de Clippel, Geoffroy and Kareen Rozen (2021). “Communication, Perception and Strategic Obfuscation”. *Working Paper*.

Dean, Mark and Nathaniel Neligh (2019). “Experimental Tests of Rational Inattention”. *Working Paper*.

Denti, Tommaso, Massimo Marinacci, and Aldo Rustichini (2022). “Experimental Cost of Information”. *American Economic Review* 112.9, pp. 3106–3123.

Dewan, Ambuj and Nathaniel Neligh (2020). “Estimating Information Cost Functions in Models of Rational Inattention”. *Journal of Economic Theory* 187.

Fischbacher, Urs (2007). “z-Tree: Zurich Toolbox for Ready-made Economic Experiments”. *Experimental Economics* 10.2.

Gentzkow, Matthew and Emir Kamenica (2014). “Costly persuasion”. *The American Economic Review* 104.5, pp. 457–462.

Heidhues, Paul and Botond Kőszegi (2018). “Behavioral industrial organization”. *Handbook of Behavioral Economics: Applications and Foundations 1* 1, pp. 517–612.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin (2021a). “Complex Disclosure”. *Management Science* 68.5, pp. 3236–3261.

— (2021b). “Is no news (perceived as) bad news? An experimental investigation of information disclosure”. *American Economic Journal: Microeconomics* 13.2, pp. 141–173.

— (2022). “Complex disclosure”. *Management Science* 68.5, pp. 3236–3261.

Lipnowski, Elliot, Laurent Mathevet, and Dong Wei (2020). “Attention management”. *American Economic Review: Insights* 2.1, pp. 17–32.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt (2023). “Rational inattention: A review”. *Journal of Economic Literature* 61.1, pp. 226–273.

Martin, Daniel (2016). “Rational inattention in games: experimental evidence”. *Available at SSRN 2674224*.

— (2017). “Strategic Pricing with Rational Inattention to Quality”. *Games and Economic Behavior*.

Martin, Daniel and Edwin Muñoz-Rodriguez (2022). “Misperceiving mechanisms: Imperfect perception and the failure to recognize dominant strategies”. *European Economic Review*.

Matějka, Filip (2015). “Rigid pricing and rationally inattentive consumer”. *Journal of Economic Theory* 158, pp. 656–678.

Matyskova, Ludmila and Alfonso Montes (2023). “Bayesian persuasion with costly information acquisition”. *Journal of Economic Theory*, p. 105678.

Mitzkewitz, Michael and Rosemarie Nagel (1993). “Experimental Results on Ultimatum Games with Incomplete Information”. *International Journal of Game Theory volume* 22, pp. 171–198.

Pomatto, Luciano, Philipp Strack, and Omer Tamuz (2023). “The Cost of Information: The Case of Constant Marginal Costs”. *American Economic Review* 113.5, pp. 1360–1393.

Rapoport, Amnon (1999). “Game Theory: Contributions to the Study of Human Cognition”. *Cognitive Studies* 6.2, pp. 142–167.

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, pp. 377–394.

Rapoport, Amnon, James Sundali, and Darryl Seale (1996). “Ultimatums in Two-person Bargaining with One-sided Uncertainty: Demand Games”. *Journal of Economic Behavior Organization* 30.2, pp. 173–196.

Ravid, Doron (2020). “Ultimatum Bargaining with Rational Inattention”. *American Economic Review* 110.9, pp. 2948–63.

Spurlino, Eric (2022). “Rationally Inattentive and Strategically (Un)sophisticated: Theory and Experiment”. *Working paper.*

Weg, Eythan, Rami Zwick, and Amnon Rapoport (1999). “An Outside Option Experiment”. *Games and Economic Behavior* 14, pp. 150–286.

Yang, Ming (2015). “Coordination with flexible information acquisition”. *Journal of Economic Theory* 158, pp. 721–738.

— (2020). “Optimality of Debt under Flexible Information Acquisition”. *The Review of Economic Studies,* 87.1, pp. 487–536.

A Experiment Implementation

A.1 Experiment Instructions (LOO Treatment)

Introduction:
You are about to participate in an experiment on the economics of decision making. You will be paid for your participation at the end of the session. The exact amount earned depends partly on chance and partly on following instructions carefully and making good decisions. You may earn a considerable amount of money in addition to the €8 participation fee.

The Game:
In this experiment you will have the opportunity to earn Experimental Points (EPs) by playing several rounds of a "buyer-seller" game. Experimental Points increase your chances of receiving real euro rewards at the end of the experiment. We will start by describing the game:
In each round, a Seller and a Buyer are matched together, and the Seller is assigned a hypothetical product that has a value of either 100 or 50 Experimental Points to the Buyer. The Seller decides whether to offer a price of 50 or 25 Experimental Points for the product, which is how many Experimental Points it will cost the Buyer to obtain the product. The Buyer is told this price, and if he or she so chooses, the Buyer has a way to learn the value before deciding whether or not to accept the Seller's offer.
If the Buyer accepts the Seller's offer, then the Buyer gets an amount equal to the value of the product minus its price and the Seller gets an amount equal its price. If the Buyer rejects the offer, the Buyer gets a fixed amount of 12.5 Experimental Points and the Seller gets 0 Experimental Points.

Before the game starts, everyone will have the chance to play a few practice rounds to familiarize with the interface and the game rules.

Game Dynamics:
The session will be divided into two phases: (1) a Seller phase and then (2) a Buyer phase. Decisions made during the Seller phase will be used to determine prices in the Buyer phase.

<u>1. Seller Phase:</u> Every subject starts out as a Seller. As a Seller, you will be asked to choose your pricing strategy for the remainder of the game. By default, all products with a value of 100 will automatically receive a price of 50. You will have to indicate with what probability (from 0% to 100%) you want to set a price of 50 (instead of a price of 25) when the product has a value of 50.

<u>2. Buyer Phase:</u> After the Seller phase, every subject will become a Buyer. In the Buyer

phase, you will complete 16 rounds. In each round you will be matched with a randomly drawn Seller, and it is likely you will be matched with a different Seller than in the preceding round.
Buyer round description:
Each round, the computer randomly selects the value of the product: With 50%chance the product has a value of 100 and with 50% chance a value of 50.

- If the value drawn is 100: A price of 50 will automatically be chosen.
- If the value drawn is 50: The Seller strategy is implemented. That is, the computer selects a high price of 50 with the probability that the Seller you were matched with chose in their Seller phase.

The Buyer is informed of the price and can learn the value of the product by performing a task of adding up 20 numbers. The 20 numbers are determined by randomly drawing twenty numbers between -50 and 50 such that they add up to the product's value. The Buyer cannot use scratch paper or a calculator. If the Buyer rejects the Seller's offer, the Buyer's payoff is 12.5, and the Seller's payoff is 0. The Buyer has up to 120 seconds to make a decision. If no decision is made after 120 seconds, the offer will be rejected.

Round Summary:
1. Randomly match every Buyer with a Seller
2. The value of the product is randomly determined to be 100 or 50.
3. Price is chosen
   a) **If the Value is 100:** the price is set to 50 by default.
   b) **If the value is 50:** Price implemented according to the Seller strategy.
4. The Buyer is told the price of the product.
5. The Buyer can learn the value of the product by adding up 20 numbers.
6. The Buyer accepts or rejects the offer.
   a) **Payoffs for accepting:** Buyer gets value minus price, Seller gets price.
   b) **Payoffs for rejecting:** Buyer gets 12.5, Seller gets 0.

Throughout the game, you will have available a payoff table for reference. This table will summarize all the possible payoff combinations derived from different values, prices, and Buyer actions.

Payment:
In addition to the €8 show-up fee, you can earn two additional rewards of €20:

The chances of winning the first €20 reward depend solely on your performance in the Seller role. We will select a random round and you will get the Experimental Points earned in that round as a Seller, which will determine your chances of earning the first reward.

The chances of winning the second €20 reward depend solely on your performance in the Buyer role. We will select another random round and you will get the Experimental Points earned in that round as a Buyer, which will determine your chances of earning the second reward.

For every Experimental Point you receive your chance of receiving that reward will increase by 1 percentage point.

- For example, imagine that the computer randomly selected the 15th round for the seller role, and the 4th round for the buyer role. Let suppose that in the 15th round as a seller you earned 50 Experimental Points, and in the 4th round as a buyer you earned 25 Experimental Points. You would then have a 50% chance of winning the first €20 reward and a 25% chance of winning the second €20 reward. Both rewards are independent of each other and your chances of earning one reward would not affect your chances in the other one.

Remember: All identities remain anonymous. No one will learn what strategy you played or what payoff you earned.

A.2 Interface Screenshots

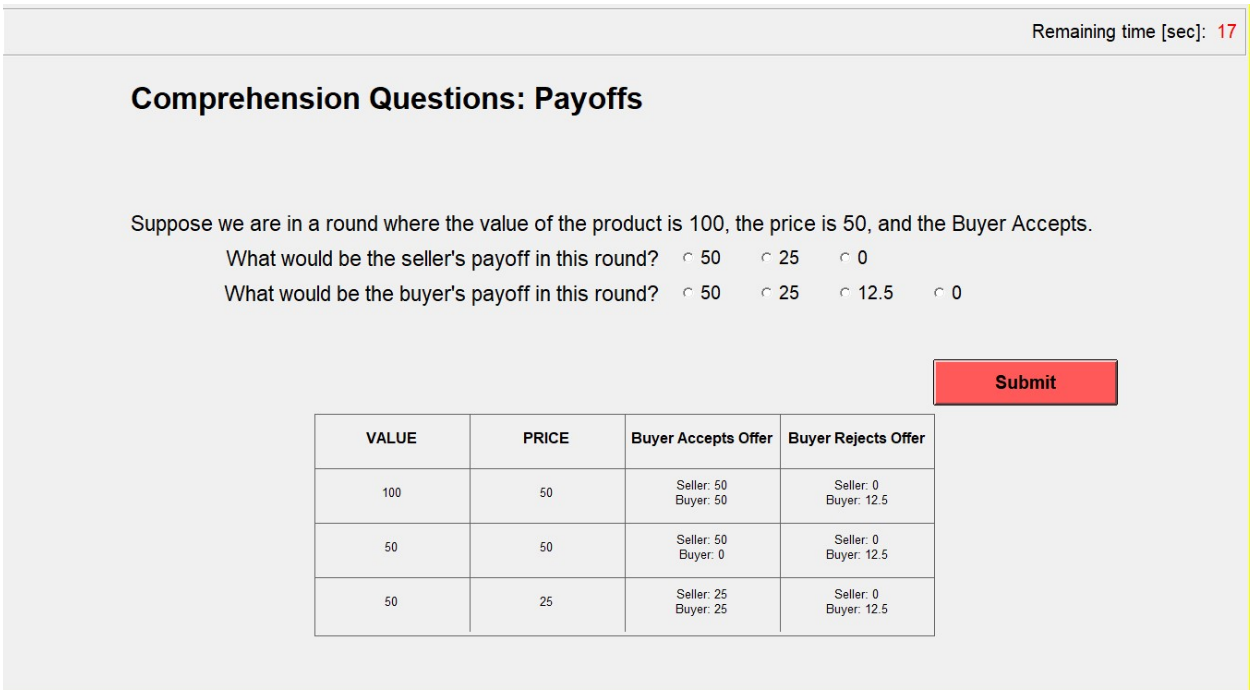

Figure 7: Comprehension question example screenshot.

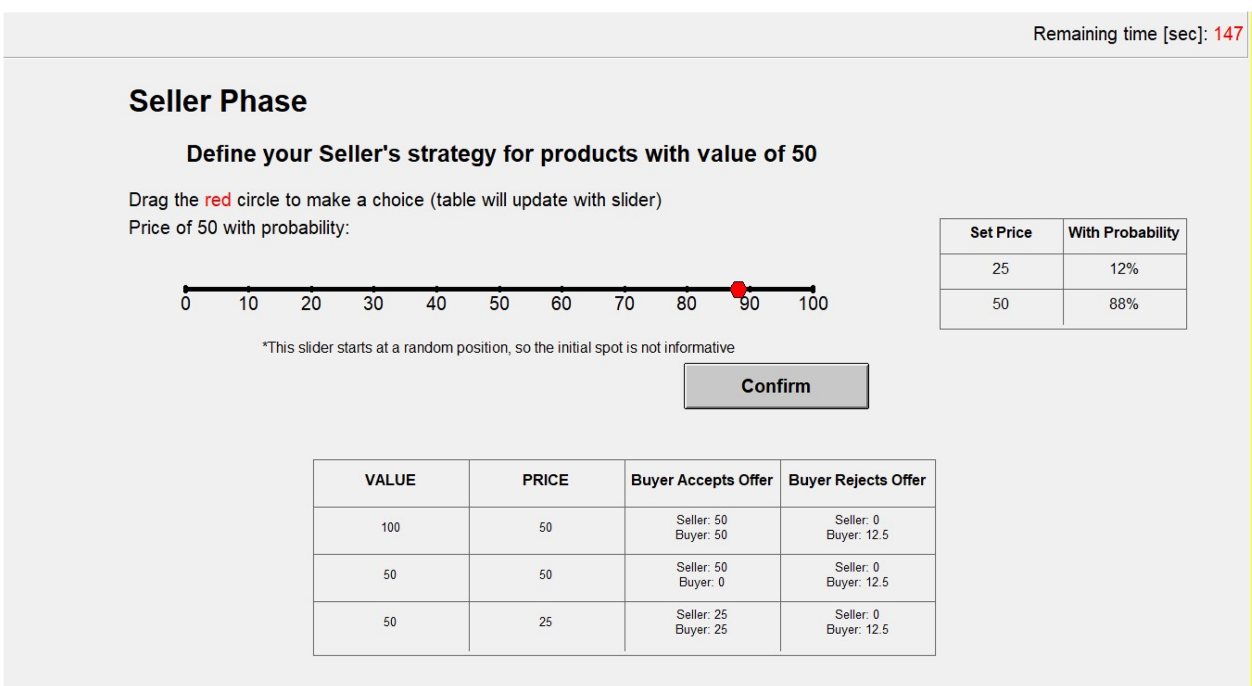

Figure 8: Seller phase example screenshot.

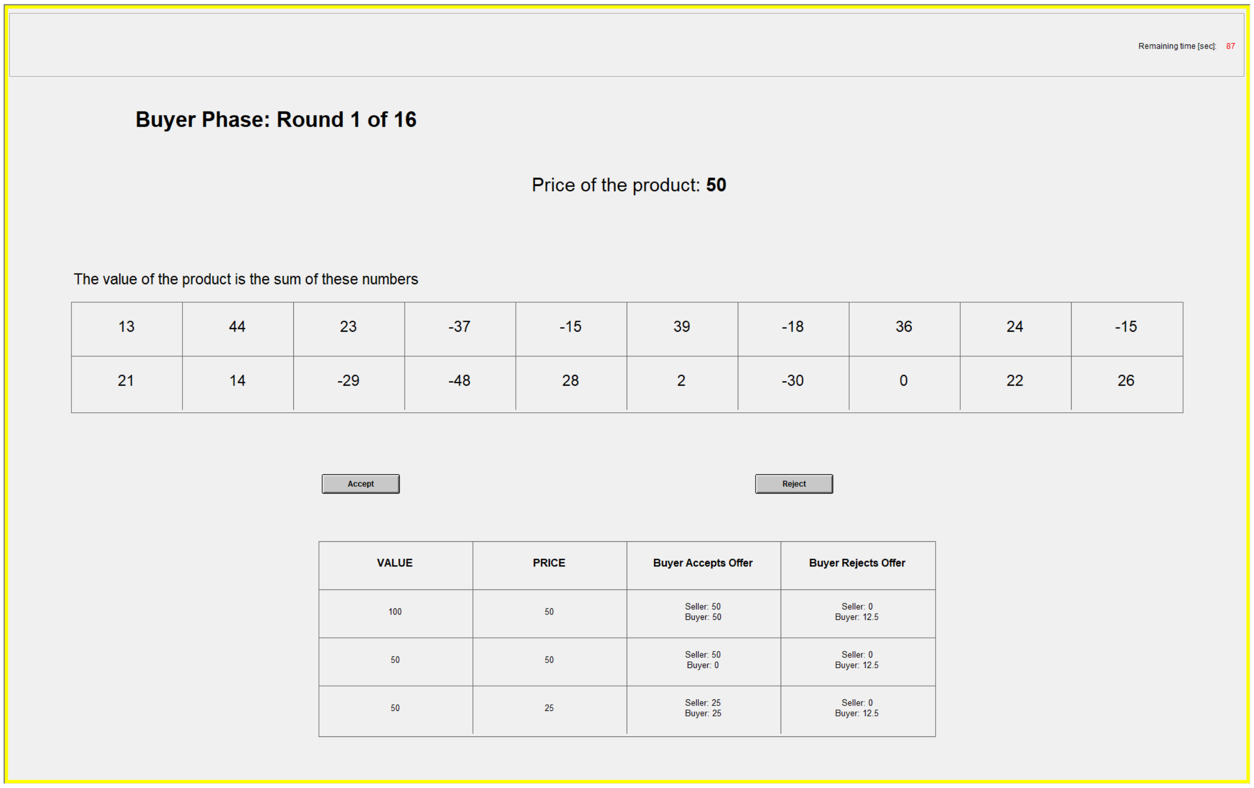

Figure 9: Buyer phase example screenshot.

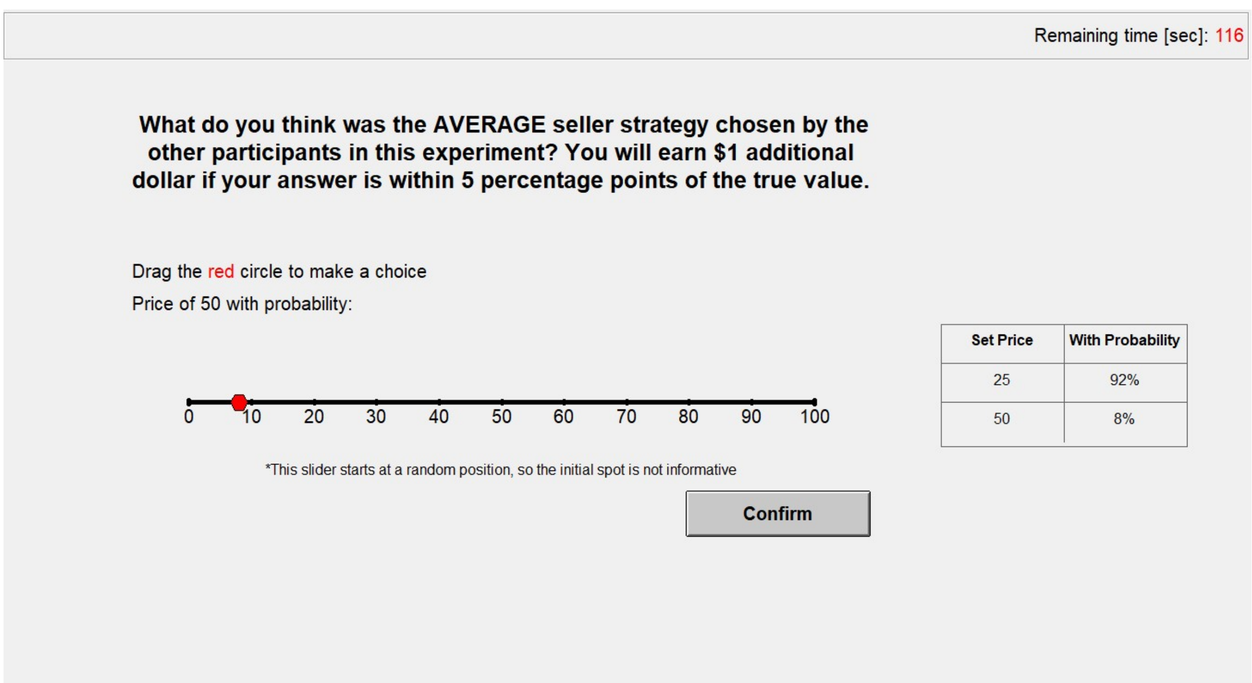

Figure 10: Seller strategy belief elicitation example screenshot.

B Additional Figures

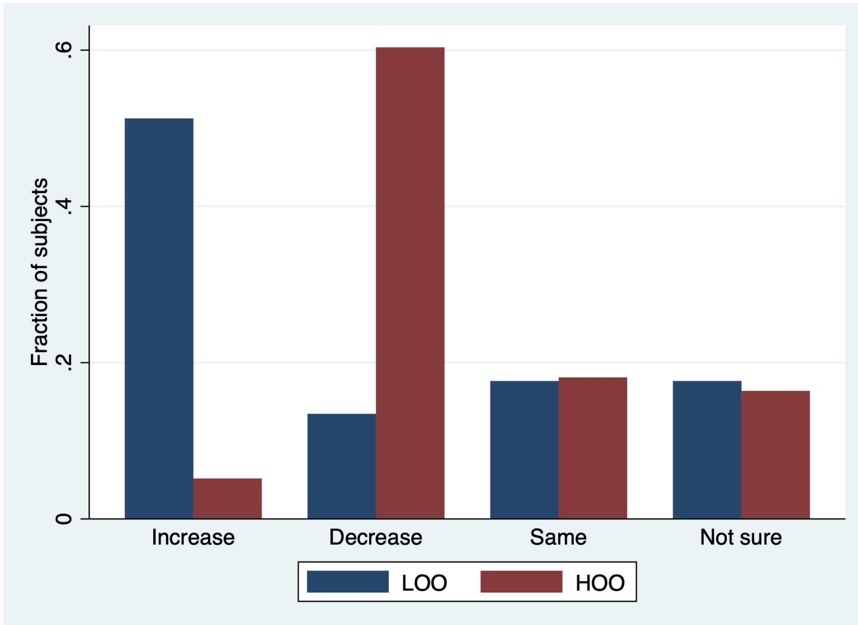

Figure 11: Self-reported change in seller strategy for hypothetical switch in outside option.

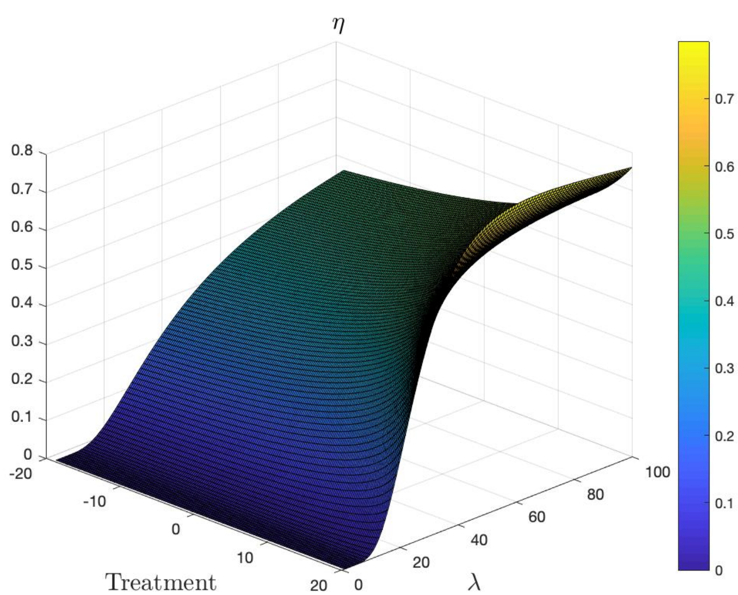

Figure 12: Theoretical relationship between mimicking, marginal cost of attention, and outside option (adjusted to account for a disutility level of $\phi = 50$).

C Additional Tables

	Rounds 1-8			Rounds 9-16			Overall		
(Value,Price)	50,25	50,50	100,50	50,25	50,50	100,50	50,25	50,50	100,50
LOO	6.50%	23.53%	24.89%	9.79%	17.62%	24.53%	8.17%	20.77%	24.71%
HOO	6.51%	23.66%	19.20%	9.36%	13.54%	19.18%	7.89%	18.32%	19.19%

Table 4: Mistake rates at each price and value for each treatment (including buyer rounds that ran out of time and hence ended up as rejections).