



Rational inattention in games: experimental evidence

David Almog¹ · Daniel Martin^{1,2} 

Received: 20 November 2023 / Revised: 27 August 2024 / Accepted: 2 September 2024 /
Published online: 13 September 2024
© The Author(s) 2024

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 consistent with this behavioral pattern.

Keywords Rational inattention · Limited attention · Cognitive costs · Strategic pricing · Asymmetric information

JEL Classification C91 · D82 · D83 · D90 · D91

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 purchase without

✉ Daniel Martin
daniel@martinonline.org

David Almog
david.almog@kellogg.northwestern.edu

¹ Kellogg School of Management, Northwestern University, Evanston, USA

² University of California, Santa Barbara, USA

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 information 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, we study 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 the buyer rejects their offer, which we exogenously vary between subjects. In real markets, variation in the outside option could be due to 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. 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 that is priced high and less likely to accept a low-value product that is priced high.

These changes in both mistakes and beliefs suggest that buyers are indeed adjusting their attention to available information in response to seller strategies. However, determining if buyers are *optimally* adjusting their attention in response to seller strategies requires modeling and measuring the costs and benefits of attention. As a starting point, we examine whether buyer behavior in our experiment is consistent with a general model of attention in which the benefits of attention are just

¹ Questions of this nature call for integrating human cognition into game theory. Professor Amnon Rapoport, in whose memory this special issue is dedicated, was an early advocate for incorporating 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 et al. (1989) provide evidence on how a bargainer can use her outside option to gain leverage, and Weg et al. (1999) build on this by showing how manipulation of the environment through outside opportunities can induce subjects to approximate rational behavior.

the monetary payoffs of playing the game. We find that the *No Improving Action Switches* (NIAS) condition of Caplin and Martin (2015), which characterizes our model of attention, is not satisfied if we take game payoffs at face value. This is because buyers reject offers at a very high rate, so much so that there would be an expected utility gain from a wholesale switch from rejecting to accepting offers. However, this finding is consistent with a behavioral literature (e.g., Mitzkewitz and Nagel (1993); Camerer and Thaler (1995); Rapoport et al. (1996)) that demonstrates individuals reject apparently unfair offers in the Ultimatum Game when there is outcome uncertainty. We move forward by amending our model of attention 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 accepting and rejecting offers optimally given their attention in each treatment.

Determining if buyers are optimally adjusting their attention in response to seller strategies also requires modeling and measuring the cognitive costs of attention. We begin by considering a leading model of rational inattention (RI) in which costs scale linearly with Shannon mutual information (Matejka, 2015).³ We derive three testable predictions from the Shannon model. First, we predict that sellers in the HOO treatment will be more likely to set a high price for low-value products. Intuitively, the loss incurred from having an offer rejected is smaller with a higher outside option, which incentivizes sellers to gamble on pricing high more often. Second, the model predicts that buyers in the HOO treatment will be less likely to accept low-value products when the price is high. Intuitively, the higher likelihood of encountering a low-value product makes buyers more suspicious when the price is high and leads them to reject more often. While this change in buyer behavior should dampen a low-value seller's propensity to price high, it is not enough to offset the boost provided by a higher outside option. We find that the first two predictions of the Shannon model are consistent with our experimental results.⁴

The third testable prediction of the Shannon model is that buyers will be *more* likely to make the mistake of rejecting high-value, high-price offers in the HOO treatment. This surprising third prediction arises because the optimal posterior belief at which buyers reject offers should not change across treatments under the Shannon model. Thus, the stronger incentive to reject high-price offers, which arises from the increase in the prevalence of low-value products, leads to an increase in the mistake of rejecting high-value products when the price is high. Instead, we find a higher probability of low value when buyers reject offers in our experiment, which suggests that the posteriors at which buyers reject offers change across treatments. This leads to a rejection of the third prediction of the Shannon model: buyers in our experiment

³ We follow Dean and Neligh (2023) 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 We generate both full-equilibrium predictions, which are valid at the individual level, and partial-equilibrium predictions, which are valid at the aggregate level, but all predictions are checked at the aggregate level given limited observations at the individual level.

are less likely to make the mistake of rejecting high-value products when the price is high.

This result offers a larger insight into 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 the invariance of posterior beliefs that appears to be violated in our experiment. On the other hand, optimal posterior beliefs can vary with the prior in RI models based on costly experiments (e.g., Denti et al. (2022); Pomatto et al. (2023)).⁵ In this second class of models, the costs of information are what is 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 the game we implement experimentally. Indeed, buyer choices in our experiment pass a version of the *No Improving Attention Cycles (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 buyer behavior in our experiment is consistent with an RI model based on costly experiments.

The rest of the paper proceeds as follows. In Sect. 1.1, we review related literature on modeling and testing inattention. In Sect. 2, we describe our experimental design and explain various design choices. In Sect. 3, we examine treatment effects in terms of seller strategies, buyer beliefs, and buyer mistakes. We study consistency with a general model of attention in Sect. 4 and with both the Shannon model and RI with costly experiments in Sect. 5. Finally, we conclude with a brief summary in Sect. 6.

1.1 Related literature

A large literature in economics, psychology, and neuroscience has provided evidence that cognitive limitations influence individual decision making. 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.⁶ Recently, a number of papers have considered the game theoretic implications of this modeling approach (e.g., Gentzkow and Kamenica (2014); Matejka (2015); Yang (2015); Martin (2017); Bloedel and Segal (2018); Boyaci and Akçay (2018); Lipnowski et al. (2020); Ravid (2020); Yang (2020); Albrecht and Whitmeyer (2023); Matyskova and Montes (2023)).

Costly information acquisition in strategic settings has been studied in a parallel line of research where informational costs are induced (e.g., Elbittar et al. (2020); Szkup and Trevino (2020)). However, most experimental results on RI are focused

⁵ In the language of Denti et al. (2022), “statistical experiments” map states of 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.

⁶ See Caplin (2016) and Mackowiak et al. (2023) for reviews of the RI literature, and Heidhues and Koszegi (2018) for a review of important IO applications.

on non-induced costs in non-strategic settings. Using individual decision-making tasks, Dean and Neligh (2023) demonstrate in a robust way that subjects adjust their attention in response to changes in rewards.⁷ 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, the subjects in their experiment largely satisfy the foundational Locally Invariant Posteriors 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 we set identify using NIAS is particular to games, and unsurprisingly, NIAS passes in their experiment without the need to accommodate additional behavioral forces. Because of this, our research introduces one of the first rational inattention models that incorporate the influence of behavioral factors on attention allocation. Recent works by Bolte and Raymond (2023) and Almog et al. (2024) suggest the importance of expanding rational attention modeling to account for emotional states.

In another non-strategic experiment, Ambuehl et al. (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. Altmann et al. (2022) find in a non-strategic setting that while attention-grabbing policies increase subjects’ focus in a choice domain, they can backfire by decreasing decision quality in other choice domains.

A related branch of the experimental literature is a growing set of papers that consider inattention to game payoffs. Brocas et al. (2014) make use of an innovative mouse-tracking 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 Schotter (2020) study the allocation of attention across games and find that subjects choose to allocate their attention to games with certain properties, like higher maximum payoffs. In a subsequent paper, Avoyan et al. (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-Rodríguez (2022) study inattention to payoffs in the Becker-DeGroot-Marschak (BDM) mechanism (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

⁷ Dewan and Neligh (2020) and Caplin et al. (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.

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 obfuscate strategically.⁸ 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. In de Clippel and Rozen (2021), the sender should aim to obfuscate always in the opposing-interest state regardless of the treatment, and the receiver's prior should change across treatments because the communication channel becomes noisier. In our experiment, the sender's decision on how to price low-value products varies with the treatment, so changes in the buyer's prior rise 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 cannot 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).⁹ Each value has an equal probability of being realized, and this probability is known to both the buyer and the seller. The seller is shown the value of the product and then offers the buyer a price for that product of either 50 or 25 EPs. The buyer is told the product's price, and if they choose to, the buyer has a way to learn the product's value by exerting

⁸ Related to this, Jin et al. (2022) show senders elect to shroud information by sending complex messages when the state is not favorable for them.

⁹ Subjects were told that EPs represented the probability of winning a €20 prize in a randomly selected round.

cognitive effort before deciding whether to accept or reject the seller's offer. If the buyer accepts the seller's offer, the buyer receives an amount equal to the value of the product minus its price, and the seller receives 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 receives $\tau \in \{0, 20\}$ 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's side unchanged. In the low outside option (LOO) treatment, sellers get $\tau = 0$ EPs if their offer is rejected. In the high outside option (HOO) treatment, the payoff of a seller whose offer is rejected increases to $\tau = 20$ EPs.

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.¹⁰ 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.¹¹ 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." This rate measures how often low-value sellers mimic the pricing 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 subject's 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.¹² 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.¹³ A round of the buyer stage ended if the buyer accepted or rejected the subsequent offer within 120 s. If no decision was made in 120 s, the offer was automatically rejected.

¹⁰ See Figure A.1 in the Appendix for an example screenshot.

¹¹ The slider started at a random location to control issues related to anchoring or default effects.

¹² See Figure A.2 in the Appendix for an example screenshot.

¹³ Subjects were not allowed to use scratch paper or a calculator.

2.2 Lab implementation and extra tasks

Our pre-registered experiment was run in person at the WISO Experimental Lab of Hamburg University.¹⁴ A total of 12 sessions were conducted, and each session was composed of between 10 and 30 subjects, for a total of 238 subjects.¹⁵ The experiment was organized and recruited using the hroot software package (Bock et al., 2014).

At the beginning of each session, subjects received a copy of the instructions, which were also read aloud before the experiment began.¹⁶ 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).¹⁷ We provided the payoff table during these questions because it 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 uniformly at 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 11.3% did not submit an answer. If subjects 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 section in which subjects were given the opportunity to practice both as a seller and a buyer without any payoff consequences. In this section, subjects first entered a practice seller strategy and then completed four buyer practice rounds. There were two main differences between the buyer practice rounds and the regular buyer rounds, which we made clear to the 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 section 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

¹⁴ Preregistration plan available at <https://aspredicted.org/wi5g4.pdf>

¹⁵ 120 and 118 subjects in the LOO and HOO treatments, respectively.

¹⁶ The instructions for the LOO treatment, which are nearly identical to the instructions for the HOO treatment, can be found in Appendix A.

¹⁷ See Figure A.3 in the Appendix for an example screenshot.

(after round 16), we elicited each subject's beliefs about the average seller strategy in their session.¹⁸

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 six empty spaces for each table without access to the actual payoff table. One purpose of this memory test was to determine whether subjects 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 et al. (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 five 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 skills separately from strategic decision making. For this task, we used 20 numbers from -100 to 100 (instead of 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 min.¹⁹ The experiment was programmed in z-Tree (Fischbacher, 2007).

¹⁸ To increase familiarity, we made the interface for belief elicitation similar to the interface from the seller stage.

¹⁹ Payoffs averaged €21.2 in the LOO treatment and €26.5 in the HOO treatment, and ranged from €8 to €56.7.

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 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 buyer's 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 are potentially in flux given 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 ample time to sum the numbers in the buyer stage (120 s) in case fatigue slowed down some subjects. In fact, we found that performance actually improved in the second half, 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 rounds. Thus, we felt that automatically setting a high price for these products would not result in 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 reasons for this design feature were 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 (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).²⁰

3.1 Seller strategies

Sellers chose a higher probability of pricing high (price = 50) when the value was low (value = 50) 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 0.0054.²¹ Thus, increasing the outside option from 0 to 20 EPs resulted in a 27% increase in the average mimicking rate. Figure 1 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 a within-subject measure of the treatment effect on seller strategies. After all extra tasks had been completed but right before the results of the game were revealed, we asked subjects the following unincientivized 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

²⁰ Including a control for whether the subject has a friend in the room is also used in Jin et al. (2021) and Jin et al. (2022).

²¹ These are the chosen mimicking rates and not the actual mimicking rates realized in the experiment, which are of a similar magnitude.

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 B.1 in the Appendix shows that a vast majority of subjects in the LOO treatment reported that they would increase their seller strategy (price high more often when the value is low) 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 on 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, some share of this effect may be due to 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 when controlling for own seller strategy (the coefficient on HOO is 4.2 and has a p -value of 0.135). Given the dynamic change in beliefs, we will focus our analysis of buyer behavior on the second half of buyer stage rounds.

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 of strategic reasoning. In the second half of rounds, the rates of rejecting low-price offers are 9.2% and 8.5%

Table 1 Subject characteristics and balance across treatments

	LOO		HOO		Comparison	
	Mean (sd.)	Obs.	Mean (sd.)	Obs.	Diff	t-stat
Male	0.401 (0.49)	112	0.5 (0.50)	112	- 0.09	1.4
Age (years)	26.4 (5.9)	120	25.9 (5.8)	118	0.55	0.7
English level (1–5)	4.04 (0.84)	117	4.05 (0.91)	118	- 0.01	0.07
Friend in session	0.116 (0.35)	120	0.042 (0.20)	118	0.074	2.1**
≥ 10/12 on payoff comp	0.891 (.31)	120	0.881 (.32)	118	0.010	0.24
12/12 on payoff comp	0.733 (0.44)	120	0.733 (0.44)	118	0	0.06
Mixed strategy comp	0.741 (0.40)	120	0.762 (0.42)	118	- 0.02	0.37
8/8 on memory task	0.633 (0.48)	120	0.703 (0.45)	118	- 0.07	1.4
Math test correct (0–5)	0.416 (0.90)	120	0.677 (1.17)	118	- 0.26	1.92*
Math test close (0–5)	0.9 (1.12)	120	1.34 (1.40)	118	- 0.44	2.7***
≥ 1/5 math test	0.258 (0.43)	120	0.338 (0.47)	118	- 0.080	1.35
≥ 1/5 math test close	0.533 (0.50)	120	0.618 (0.48)	118	- 0.08	1.3

Observation is per subject. Value is missing if demographic information was not provided by the subject. All variables without a specified type in parentheses are dummies * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

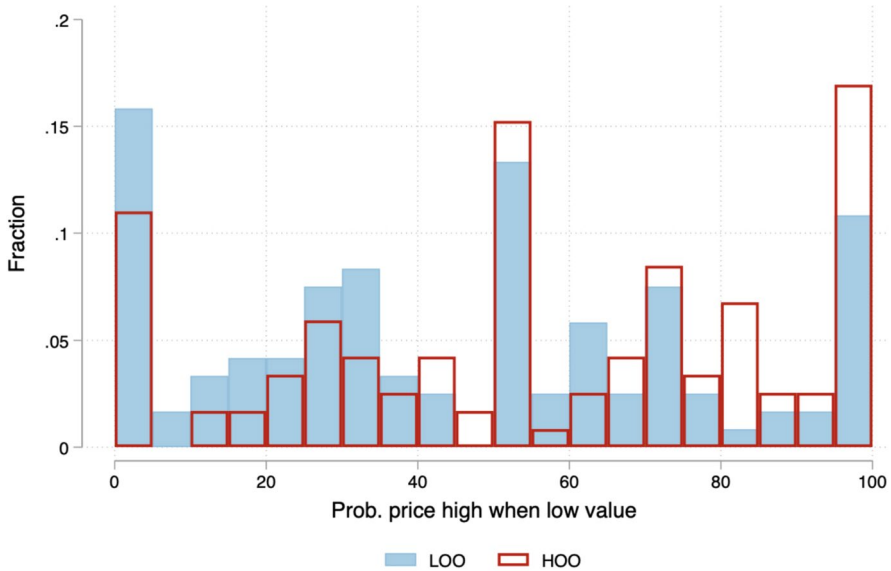


Fig. 1 Distribution of seller strategies (mimicking rates) by treatment

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,

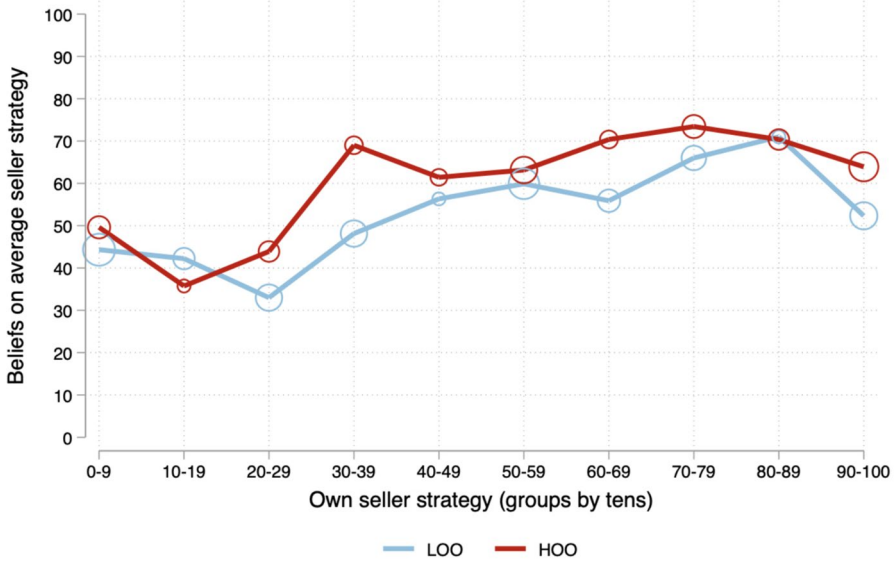


Fig. 2 Treatment effect on elicited beliefs contingent on own seller strategy. Bubbles provide the frequency of each seller strategy bin. Elicited beliefs are from the end of the buyer stage

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 this improvement does not rule out fatigue entirely; 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 a high-value product, and subjects in the HOO treatment had a lower mistake rate of rejecting such offers: 17.8% versus 23.4% (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

Table 2 Mistake rates at each price and value by treatment

(Value, price)	Rounds 1–8 (%)			Rounds 9–16 (%)			Overall (%)		
	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

had a lower mistake rate of accepting such offers: 13.7% versus 18.2% (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 an 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 treatment dummy variable, as well as dummy variable controls for the number of correct questions on the math test and 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 (compared to the 21.4% pooled mistake rate of LOO). This result is significant at a 5% level without clustering, and when we cluster standard errors at a subject level, it is significant at a 10% level. The two controls we added appear to be relevant, as individuals who performed better on the math test and those who participated with a friend in the experiment exhibited significantly lower 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.

4 The benefits of attention

The changes in buyer beliefs and mistakes we observe across treatments suggest buyers are indeed adjusting their attention to available information in response to seller strategies. However, to determine if buyers are *optimally* adjusting their attention in response to seller strategies requires modeling and measuring the costs and benefits of attention. As a starting point, we examine whether buyer behavior is consistent with the general model of attention characterized by Caplin and Martin (2015). This is a useful starting point because it reveals that the benefits of attention are not just monetary. Instead, subjects appear to be impacted by a previously identified psychological cost.

4.1 Model of Buyer attention

As is standard in models of attention, we treat the buyer as a decision-maker who starts off with prior beliefs about the state of the world, receives informative signals about the state, updates their beliefs, and chooses the actions that maximize expected utility given their posterior beliefs. In our setting, the state $\theta \in \Theta = \{\theta_L, \theta_H\}$ represents the product's value. The high state (θ_H) is realized with probability δ and the low state (θ_L) with probability $1 - \delta$. In our experiment, these parameters took a value of $\theta_H = 100$, $\theta_L = 50$, and $\delta = .5$.

When the buyer accepts the seller's offer, they get the product's value θ minus the price p , which can either be high (p_H) or low (p_L). If the buyer rejects the seller's offer,

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

	(1)	(2)	(3)	(4)
	Mistake	Mistake	Mistake	Mistake
High outside option	-0.0528** (0.0212)	-0.0420** (0.0213)	-0.0474** (0.0214)	-0.0474* (0.0285)
Math test correct = 1		-0.141*** (0.0291)	-0.134*** (0.0293)	-0.134*** (0.0373)
Math test correct = 2		-0.144*** (0.0378)	-0.149*** (0.0378)	-0.149*** (0.0340)
Math test correct = 3		-0.0987 (0.0804)	-0.102 (0.0804)	-0.102 (0.0724)
Math test correct = 4		-0.0952 (0.0732)	-0.0994 (0.0731)	-0.0994 (0.0721)
Math test correct = 5		-0.119 (0.0782)	-0.125 (0.0782)	-0.125 (0.101)
Friend			-0.0759* (0.0398)	-0.0759* (0.0399)
Cluster	None	None	None	Subject
N	1415	1415	1415	1415

Standard errors in parentheses * $p < .10$, ** $p < .05$, *** $p < .01$

they get their outside option β . In our experiment, prices were $p_H = 50$ and $p_L = 25$, and the buyer’s outside option was $\beta = 12.5$. Thus, when the price was low, accepting was a dominant strategy for the buyer because $\theta_H - p_L = 75 > \beta = 12.5$ and $\theta_L - p_L = 25 > \beta = 12.5$. On the other hand, when the price was high, the buyer would want to reject the seller’s offer if the value was low because $\theta_L - p_H = 0 < \beta = 12.5$.

We allow the buyer’s beliefs and attention to adjust with both the product’s price and the seller’s outside option. However, because there is a dominant action when prices are low, we focus on the buyer’s beliefs and attention when prices are high. At a given seller outside option, the buyer first forms prior beliefs μ about the probability of the product’s value. Since sellers always price high when the value is high, the buyer’s prior belief that the product’s value is high at a high price is given by the seller’s mimicking rate η (the rate at which they price high when the value is low) and the unconditional probability δ of high value:

$$\mu(\theta_H) = \frac{\delta}{\delta + (1 - \delta)\eta} \tag{1}$$

After receiving a signal of the state θ , the buyer forms posterior $\gamma \in \Gamma$, where $\gamma(\theta)$ is the probability of state θ in Θ . The buyer’s posterior that the product’s value is high ($\theta = \theta_H$) is given by

$$\gamma(\theta_H) = \frac{\mu(\theta_H)\pi(\gamma|\theta_H)}{\sum_{\theta \in \Theta} \mu(\theta)\pi(\gamma|\theta)} \tag{2}$$

Finally, 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.

The buyer's attention is represented as an information structure π , which stochastically generates posterior beliefs satisfying Eq. 2. Let $\pi(\gamma)$ be the unconditional probability of posterior $\gamma \in \Gamma$ and $\pi(\gamma|\theta)$ be the probability of posterior γ given state θ . Further, let $\Gamma(\pi) \subset \Gamma$ denote the support of a given π , which we assume to be finite for notational ease.²²

The results of Caplin and Martin (2015) can be applied to show that buyer choices are consistent with this model of attention if and only if they satisfy the NIAS condition for both treatments. The NIAS condition requires that there should be no utility gain from wholesale action switches. 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 (at a given price and seller outside option) and $P(\theta|accept)$ as the probability of state θ when buyers accept (at a given price and seller outside option),²³ this condition can be written as:

$$P(\theta_H|reject) * 12.5 + P(\theta_L|reject) * 12.5 \geq P(\theta_H|reject) * 50 + P(\theta_L|reject) * 0 \quad (3)$$

$$P(\theta_H|accept) * 50 + P(\theta_L|accept) * 0 \geq P(\theta_H|accept) * 12.5 + P(\theta_L|accept) * 12.5 \quad (4)$$

Despite the apparent weakness of this test, we find that this condition is not satisfied on aggregate for either treatment when prices are high. Thus, it does not appear that buyer choices in our experiment are consistent with a model in which the only benefits to attention are the monetary game payoffs. 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 seemingly unfair offers, even if it means receiving no return (see Mitzkewitz and Nagel (1993), Camerer and Thaler

²² Formally, this is all γ for which $\pi(\gamma|\theta) > 0$ for some $\theta \in \Theta$.

²³ Caplin and Martin (2015) refer to the joint distribution of actions and states $P(accept, \theta)$ as "state-dependent stochastic choice data".

(1995), and Rapoport et al. (1996).²⁴ 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 of 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:

$$P(\theta_H|reject) * 12.5 + P(\theta_L|reject) * 12.5 \geq P(\theta_H|reject) * 50 + P(\theta_L|reject) * \phi \quad (5)$$

$$P(\theta_H|accept) * 50 + P(\theta_L|accept) * \phi \geq P(\theta_H|accept) * 12.5 + P(\theta_L|accept) * 12.5 \quad (6)$$

This reduces to:

$$\frac{P(\theta_H|accept) * 50 - 12.5}{P(\theta_H|accept) - 1} \leq \phi \leq \frac{P(\theta_H|reject) * 50 - 12.5}{P(\theta_H|reject) - 1} \quad (7)$$

Looking first at state-dependent stochastic choice data for buyers in the LOO treatment, the bounds on ϕ are -387.9 to -14.7 . The state-dependent stochastic choice data for buyers in the HOO treatment produce wider bounds: -424.3 to -2.5 .

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 who did not perform well in the math test. The justification for this assumption is that since the math test is not a game, performance on the test should be orthogonal to this behavioral factor, which is related to strategic play. Looking just at those in the LOO treatment who performed poorly on the math test (had no problems correct on the math test), the estimated ϕ is tighter than before: -289.91937 to -21.25 . In the sections that follow, we will choose an arbitrary value from this range ($\phi = -50$) to conduct our analyses.

If we incorporate this disutility of accepting unfair offers into our model of buyer attention, the NIAS condition passes for high prices, which means we *can* model buyers as if they accept and reject offers optimally given the attention that they paid in each treatment. In other words, the mistake reduction we found is evidence in favor of an attentional reaction from buyers, 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.

²⁴ In other bargaining environments, such as the Tunisian Bazaar mechanism, Rapoport et al. (1995) present comparable evidence, indicating that buyers appeared to “punish” sellers by rejecting offers made “too late.” In that context, the sensation of unfairness arises from the slow decrease in prices.

5 Optimal attention adjustment?

To test whether buyers adjust their attention optimally in response to the costs and benefits of attention, 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. In this section, we use a value of $\phi = -50$ to reflect the disutility of accepting an 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

$$K(\pi, \lambda, \mu) = \lambda \left(\left[\sum_{\gamma \in \Gamma(\pi)} \pi(\gamma) \sum_{\theta \in \Theta} [\gamma(\theta) \ln(\gamma(\theta))] \right] - \sum_{\theta \in \Theta} [\mu(\theta) \ln(\mu(\theta))] \right) \tag{8}$$

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 seller outside options β and τ , respectively (buyer and seller payoffs in case the offer is rejected).

With this expansion, the results in Martin (2017) can be directly applied to show that if the buyer is rationally inattentive as in the Shannon model, there is a unique equilibrium of the buyer-seller game in our experiment in which the buyer’s information structure only generates two posteriors when prices are high: one at which the buyer always accepts the offer (γ^1) and one at which the buyer always rejects the offer (γ^0). The probability of the high state at these posteriors is given by:

$$\gamma^0(\theta_H) = \min \left\{ \frac{\exp\left(\frac{\beta}{\lambda}\right) - \exp\left(\frac{\theta_L - p_H + \phi}{\lambda}\right)}{\exp\left(\frac{\theta_H - p_H}{\lambda}\right) - \exp\left(\frac{\theta_L - p_H + \phi}{\lambda}\right)}, \mu(\theta_H) \right\} \tag{9}$$

and

$$\gamma^1(\theta_H) = \max \left\{ \frac{\exp\left(\frac{\theta_H - p_H}{\lambda}\right)}{\exp\left(\frac{\beta}{\lambda}\right)} \gamma^0(\theta_H), \mu(\theta_H) \right\} \quad (10)$$

Note that $\gamma^0(\theta_H) = \mu(\theta_H)$ or $\gamma^1(\theta_H) = \mu(\theta_H)$ only if the buyer pays no attention at all, in which case $\gamma^0(\theta_H) = \gamma^1(\theta_H) = \mu(\theta_H)$ and the same action distribution is chosen in both states. Outside of this special case, neither the prior nor the seller's outside option τ impacts either posterior.

In addition, the probability of each posterior (and hence each action) when prices are high is given by

$$\pi(\gamma^1) = \min \left\{ \frac{\mu(\theta_H) - \gamma^0(\theta_H)}{\gamma^1(\theta_H) - \gamma^0(\theta_H)}, 1 \right\} \quad (11)$$

and

$$\pi(\gamma^0) = 1 - \pi(\gamma^1) \quad (12)$$

Unlike the optimal posteriors, the mimicking rate does depend on the seller's outside option τ through its impact on the prior μ . Hence, this aspect of the optimal information structure does depend on the seller's strategy.

Finally, the seller's mimicking rate η in equilibrium is given by

$$\eta = \min \left\{ \frac{\delta}{1 - \delta \gamma^0(\theta_H)(1 - \gamma^1(\theta_H)) + \frac{p_L - \tau}{p_H - \tau} (\gamma^1(\theta_H) - \gamma^0(\theta_H))}, 1 \right\} \quad (13)$$

The mimicking rate is determined by several factors, such as the unconditional probability of the high state (δ), the buyer's optimal posteriors (γ_1 and γ_0), price levels (p_L and p_H), and the seller's outside option (τ).

Taken together, this system of equations allows us to make equilibrium predictions about buyer beliefs, unconditional action choices, mistake rates, and seller strategies. The first four equations (Eq. 9 to Eq. 12) also allow us to make partial-equilibrium predictions for buyer beliefs, unconditional action choices, and mistake rates for a given mimicking rate η .

5.2 Prediction for Seller strategies

As mentioned previously, optimal posteriors are invariant to changes in the seller's mimicking rate as long as buyers pay some attention to value. As a result, changes in the seller's outside option produce a clean comparative static in terms of the seller's equilibrium mimicking rate. Formally, introducing variation in the seller's outside option τ results in the following comparative static when $\eta < 1$:

$$\frac{\partial \eta}{\partial \tau} = \frac{\delta}{1 - \delta} \frac{(1 - \gamma^0(\theta_H))(1 - \gamma^1(\theta_H)) \frac{p_H - p_L}{(p_H - \tau)^2} (\gamma^1(\theta_H) - \gamma^0(\theta_H))}{\left(\gamma^0(\theta_H)(1 - \gamma^1(\theta_H)) + \frac{p_L - \tau}{p_H - \tau} (\gamma^1(\theta_H) - \gamma^0(\theta_H))\right)^2} > 0 \tag{14}$$

More generally, the seller (weakly) increases their mimicking rate as their outside option increases.

Intuitively, the loss incurred from having an offer rejected is smaller with a higher outside option, which incentivizes sellers to gamble on pricing high more often. Technically, this occurs because in order to make sellers indifferent between pricing high and low, demand for low-quality, high-price offers must decrease if the outside option increases, which can only happen if the mimicking rate increases.

Equation 14 motivates our first prediction, which is for sellers in our experiment:

Prediction 1: Sellers in the HOO treatment will mimic at a higher rate than sellers in the LOO treatment.

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

5.3 Predictions for Buyer mistakes

Under the assumption that $\eta < 1$, increasing the seller’s outside option τ results in the comparative static prediction that the buyer will accept offers less often when the outside option is higher:

$$\frac{\partial \pi(\gamma^1)}{\partial \tau} = \frac{1}{\gamma^1(\theta_H) - \gamma^0(\theta_H)} \frac{\partial \mu(\theta_H | p_H)}{\partial \tau} < 0 \tag{15}$$

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

Panel (a) of Fig. 3 shows how the mistake of accepting low-value and high-priced offers is predicted to vary with the marginal cost of attention (λ) in equilibrium for the parameter values in our experiment and $\phi = -50$. Intuitively, the increase in low-value products makes buyers more suspicious when the price is high and leads them to reject more often. While this change in buyer behavior dampens the low-value seller’s propensity to price high, it is not enough to offset the boost provided by a higher outside option. As a result, the mistake of accepting a low-value and high-priced offer is predicted to be lower in equilibrium for the HOO treatment whenever the mistake rate is positive in the LOO treatment. The reason why it falls to zero with λ is that in equilibrium, the mimicking rate increases with λ (for the parameter values in our experiment), so the rate of rejection goes towards one as the mimicking rate increases.

This same ordering is true if buyers best respond to the empirical distribution of seller strategies in our experiment (the average mimicking rate we observed in each treatment of our experiment), as shown in Panel (b) of Fig. 3. Because the mimicking rate is fixed in this partial equilibrium analysis, the mistake rate increases with λ because it gets harder to

discern the true state. Because this analysis takes the mimicking rate η as given, this result would also hold for any possible population distribution of cognitive costs λ .

These two results lead us to our second main prediction:

Prediction 2: The buyer's mistake rate of accepting a low-value product when prices are high is lower in the HOO treatment.

The experimental evidence and theoretical predictions for the Shannon model are also aligned with this prediction, as buyers made this type of mistake for 18.2% of low-value, high-price offers in the LOO treatment compared 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 Fig. 4 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 for the HOO treatment. The same ordering is true if buyers best respond to the empirical distribution of seller strategies in our experiment, as shown in Panel (b) of Fig. 4. Once again, the latter ordering would also hold for any possible population distribution of cognitive costs λ .

This time, the mistake rate changes with the marginal cost of attention similarly between equilibrium and the empirical best response. In both cases, the higher mimicking rate observed in HOO influences buyer behavior through the prior, decreasing their belief that a product is of high value when they observe a high price. However, the gap in mistakes between treatments is larger in the equilibrium prediction because higher rates of rejections dissuade mimicking.

These two results motivate our third main prediction:

Prediction 3: The buyer's mistake rate of rejecting a high-value product when prices are high is higher in the HOO treatment.

This surprising third prediction arises because, under the Shannon model, the optimal posterior belief at which buyers reject offers should not change across treatments. Thus, the stronger incentive to reject high-price offers, due to the increase in the prevalence of low-value products, leads to an increase in the mistake of rejecting high-value products when the price is high.

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 rejected high-price, high-value offers less often in the HOO treatment (17.8%, compared to 23.4% in the LOO treatment).

5.4 Attention as costly experiments

As noted previously, the optimal posteriors under the Shannon model (characterized in Eq. 9 and Eq. 10) are not impacted by the prior probability of the state (as long as the buyer pays some attention to the state). This local invariance of optimal posteriors to the prior under 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 in our experiment.

One implication is that optimal posteriors do not change across treatments, as the only impact of the seller's outside option on the buyer's problem is through changes in the seller's mimicking rate. Given that optimal posteriors do not change with the seller's

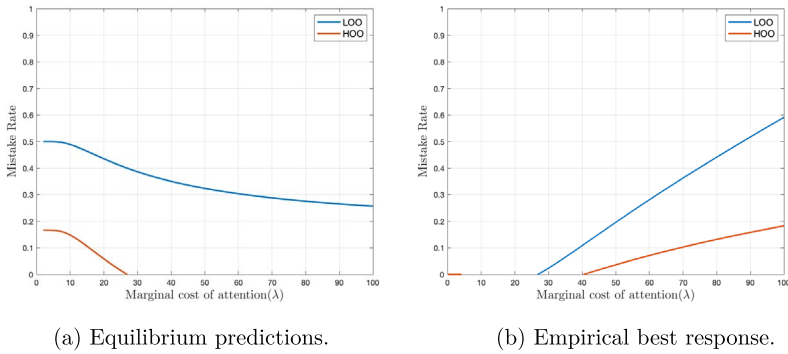


Fig. 3 Probability of accepting a low-value product at a high price for the parameter values in our experiment and $\phi = -50$

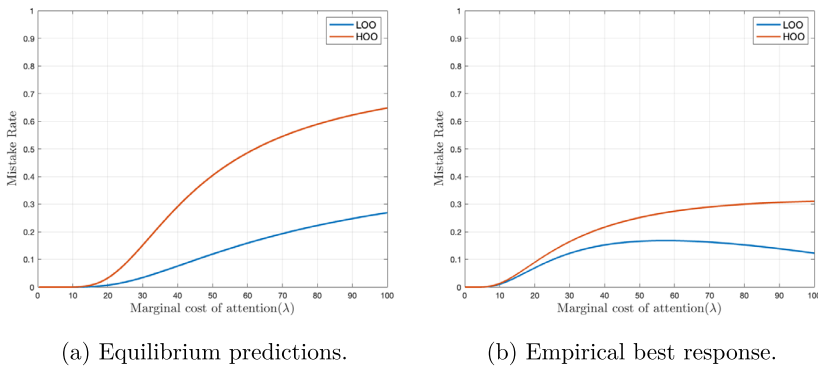


Fig. 4 Probability of rejecting a high-value product at a high price for the parameter values in our experiment and $\phi = -50$

outside option, buyers should be equally well informed when accepting and rejecting at both outside option levels.²⁵ But buyers in our experiment appear to be better informed when both accepting and rejecting offers. Figure 5 shows that the “revealed” posteriors (the average likelihood at each action) are “Blackwell-ordered” across treatments, which means that the revealed posteriors for the LOO treatment lie inside of the convex hull of the revealed posteriors for the HOO treatment.²⁶ Even though the prior probability of a high value is lower for high-priced offers in the HOO treatment, and subjects believe

²⁵ The only exception would be if subjects chose to be completely inattentive in one treatment because the seller was mimicking at a very high or very low rate in that treatment.

²⁶ The Blackwell-ordering of revealed posteriors is studied by Caplin and Martin (2021).

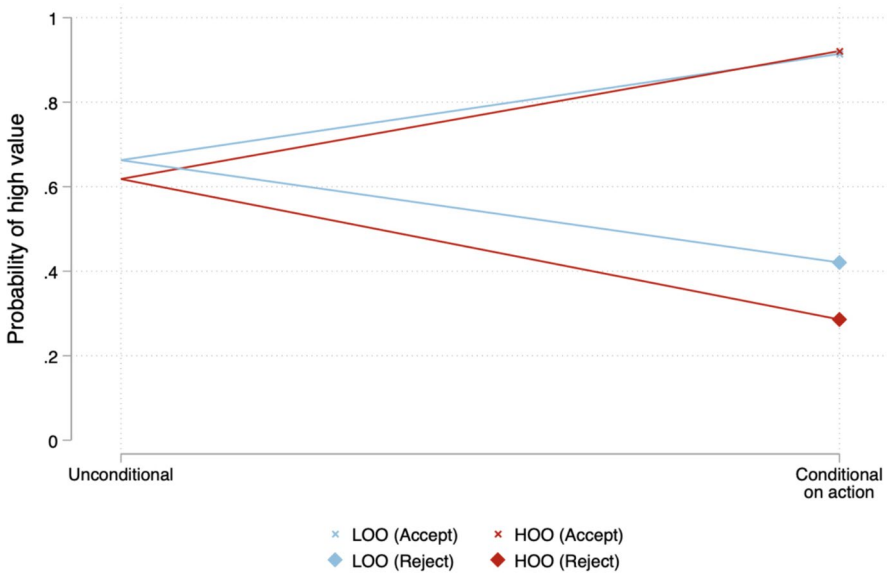


Fig. 5 Probability of high value conditional on action

this to be the case,²⁷ 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).

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 et al. (2022); Pomatto et al. (2023)). Unlike the Shannon model provided earlier, the costs in this class of models are not on the posteriors themselves. As noted in Denti et al. (2022) and Pomatto et al. (2023), this is advantageous to use in strategy settings in which one player can impact the prior over states but not the attentional costs of the opponent.

Buyer choices are consistent with this model if they pass an adapted version of the NIAC condition of 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. In Appendix D we show that this condition is satisfied.

²⁷ 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 (.618) is lower than in the LOO treatment (.663) because buyers correctly anticipate more mimicking, and observing a high price becomes a less informative signal of high value as mimicking increases.

5.5 Prior variation

As noted previously, the buyer's optimal posteriors do not change with the Shannon model, but the buyer's choice of information structure does change because the probability that buyers reach each posterior is determined in part by the seller's mimicking rate (as shown in Eqs. 11 and 12). However, the buyer's actions at each posterior do not change: the buyer always accepts at one posterior and rejects at the other.

An alternative model is that the attention of buyers does not change with the prior (i.e., their signal structure does not change), but that they change their actions at some posteriors, such as changing their cutoff belief at which they take a particular action. However, the same pattern that invalidates the Shannon model, the Blackwell ordering of revealed posteriors across treatment, also invalidates this alternative theory. If the prior $\mu(\theta_H)$ decreases with the outside option, then if the signal structure remains fixed across outside options, for each signal realization the posterior of high value will be lower with the higher outside option. No decision rule for acceptance could then simultaneously lead to a decrease in the average posterior of a high value for rejection and an increase in the average posterior of a high value for acceptance. Thus, the shift in revealed posteriors we observe across treatments also invalidates this alternative model in which the decision rule changes with the prior but not the signal structure.

6 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 is that buyers foresee this mimicking upsurge and react by increasing their attention. Buyers make fewer mistakes regardless of the value, which is, to the best of our knowledge, the first evidence of an attentional response in a strategic environment triggered solely by inference about the opponent's strategy. We find that the buyers' attentional reaction is *inconsistent* 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.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10683-024-09843-z>.

Acknowledgment We thank Anastasiia Morozova and Leshan Xu for excellent research assistance. We also thank Roberto Weber, Rami Zwick, and two anonymous reviews for many excellent suggestions and Ala Avoyan, Andrew Caplin, Geoffroy de Clippel, Michael Powell, and audiences in the Kellogg Strategy

Brown Bag Seminar, 4th Workshop on Choice-Process Data, ITAM Alumni Conference, 15th Nordic Conference on Behavioral and Experimental Economics, and WEBEAS for helpful comments
Forthcoming in *Experimental Economics*' Special Issue in memory of Amnon Rapoport.

Funding The authors have no financial or non-financial interests directly or indirectly related to the work submitted for publication. Martin's research budget provided funding for the experiments reported in this paper. The IRB ID at Northwestern University for this experiment is STU00205849. The replication material for the study is available at <https://osf.io/qayc7/>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Albrecht, Brian C , & Whitmeyer ,Mark(2023).Comparison shopping: Learning before buying from duopolists. arXiv preprint [arXiv:2302.06580](https://arxiv.org/abs/2302.06580).
- Almog, David, Gauriot,Romain, Page,Lionel, & Martin,Daniel (2024). AI Oversight and Human Mistakes: Evidence from Centre Court. arXiv preprint [arXiv:2401.16754](https://arxiv.org/abs/2401.16754).
- Altmann, Steffen, Grunewald, Andreas, & Radbruch, Jonas. (2022). Interventions and Cognitive Spillovers. *The Review of Economic Studies*, 89, 2293–2328.
- Ambuehl, Sandro, Ockenfels,Axel, & Stewart,Colin (2022). Who Opts In? Selection and Disappointment through Participation Payments. *Review of Economics and Statistics*, 1–45.
- Avoyan, Ala, Ribeiro,Mauricio, Schotter,Andrew, Schotter,Elizabeth R., Vaziri,Mehrdad, & Zou,Minghao (2023). Planned vs. Actual Attention. *Management Science*, 1–22.
- Avoyan, Ala, & Romagnoli, Giorgia. (2023). Paying for inattention. *Economics Letters*, 226, 111114.
- Avoyan, Ala & Schotter,Andrew, (2020). Attention in games: An experimental study. *European Economic Review* 124.
- Binmore, Ken, Shaked, Avner, & Sutton, John. (1989). An Outside Option Experiment. *The Quarterly Journal of Economics*, 104(4), 753–770.
- Bloedel, Alexander W , & Segal,Ilya R (2018). Persuasion with rational inattention. Available at SSRN 3164033
- Bock, Olaf, Baetge, Ingmar, & Nicklisch, Andreas. (2014). Hroot: Hamburg Registration and Organization Online Tool. *European Economic Review*, 71, 117–120.
- Bolte, Lukas , & Raymond,Collin (2023). Emotional Inattention. Working paper.
- Boyaci, Tamer, & Akçay, Yalçin. (2018). Pricing when customers have limited attention. *Management Science*, 64(7), 2995–3014.
- Brocas, Isabelle, Carrillo, Juan D., Wang, Stephanie W., & Camerer, Colin. (2014). Imperfect choice or imperfect attention? Understanding strategic thinking in private information games. *The Review of Economic Studies*, 81, 944–970.
- Camerer, Colin. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton: Princeton University Press.
- Camerer, Colin, & Thaler, Richard. (1995). Anomalies: ultimatums, dictators and manners. *Journal of Economic Perspectives*, 9(2), 209–219.
- Caplin, Andrew. (2016). Measuring and modeling attention. *Annual Review of Economics*, 8, 379–403.
- Caplin, Andrew, Csaba, Dániel., Leahy, John, & Nov, Oded. (2020). Rational inattention, competitive supply, and psychometrics. *The Quarterly Journal of Economics*, 135(3), 1681–1724.
- Caplin, Andrew, & Dean, Mark. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7), 2183–2203.
- Caplin, Andrew, & Martin,Daniel (2015). A Testable Theory of Imperfect Perception. *The Economic Journal* 125.February, 184–202.

- Caplin, Andrew, & Martin, Daniel. (2021). Comparison of decisions under unknown experiments. *Journal of Political Economy*, 129(11), 3185–3205.
- de Clippel, Geoffroy, & Rozen, Kareen (2021). Communication, Perception and Strategic Obfuscation. Working Paper.
- Dean, Mark, & Neligh, Nathaniel. (2023). Experimental tests of rational inattention. *Journal of Political Economy*, 131(12), 3415–3461.
- Denti, Tommaso, Marinacci, Massimo, & Rustichini, Aldo. (2022). Experimental cost of information. *American Economic Review*, 112(9), 3106–3123.
- Dewan, Ambuj, & Neligh, Nathaniel (2020). Estimating Information Cost Functions in Models of Rational Inattention. *Journal of Economic Theory* 187.
- Elbittar, Alexander, Gomberg, Andrei, Martinelli, César., & Palfrey, Thomas R. (2020). Ignorance and bias in collective decisions. *Journal of Economic Behavior and Organization*, 174, 332–359.
- Fischbacher, Urs. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10, 2.
- Gentzkow, Matthew, & Kamenica, Emir. (2014). Costly persuasion. *The American Economic Review*, 104(5), 457–462.
- Heidhues, Paul, & Koszegi, Botond. (2018). Behavioral industrial organization. *Handbook of Behavioral Economics: Applications and Foundations*, 1(1), 517–612.
- Jin, Ginger Zhe, Luca, Michael, & Martin, Daniel. (2021). Is no news (perceived as) bad news? An experimental investigation of information disclosure. *American Economic Journal: Microeconomics*, 13(2), 141–173.
- Jin, Ginger Zhe, Luca, Michael, & Martin, Daniel. (2022). Complex disclosure. *Management Science*, 68(5), 3236–3261.
- Lipnowski, Elliot, Mathevet, Laurent, & Wei, Dong. (2020). Attention management. *American Economic Review: Insights*, 2(1), 17–32.
- Mackowiak, Bartosz, Matejka, Filip, & Wiederholt, Mirko. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1), 226–273.
- Martin, Daniel (2016). Rational inattention in games: Experimental evidence. Available at SSRN 2674224.
- Martin, Daniel (2017). Strategic Pricing with Rational Inattention to Quality. *Games and Economic Behavior*.
- Martin, Daniel, & Muñoz-Rodríguez, Edwin (2022). Misperceiving mechanisms: Imperfect perception and the failure to recognize dominant strategies. *European Economic Review*.
- Matejka, Filip. (2015). Rigid pricing and rationally inattentive consumer. *Journal of Economic Theory*, 158, 656–678.
- Matyskova, Ludmila, & Montes, Alfonso (2023). Bayesian persuasion with costly information acquisition. *Journal of Economic Theory*, p. 105678.
- Mitzkewitz, Michael, & Nagel, Rosemarie. (1993). Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory*, 22, 171–198.
- Pomatto, Luciano, Strack, Philipp, & Tamuz, Omer. (2023). The cost of information: The case of constant marginal costs. *American Economic Review*, 113(5), 1360–1393.
- Rapoport, Amnon. (1999). Game theory: Contributions to the study of human cognition. *Cognitive Studies*, 6(2), 142–167.
- Rapoport, Amnon, Erev, Ido, & Zwick, Rami. (1995). An experimental study of Buyer/Seller negotiation with one-sided incomplete information and time discounting. *Management Science*, 41(3), 377–394.
- Rapoport, Amnon, Sundali, James, & Seale, Darryl. (1996). Ultimatums in two-person bargaining with one-sided uncertainty: Demand games. *Journal of Economic Behavior and Organization*, 30(2), 173–196.
- Ravid, Doron. (2020). Ultimatum bargaining with rational inattention. *American Economic Review*, 110(9), 2948–63.
- Spurlino, Eric (2022). Rationally Inattentive and Strategically (Un)sophisticated: Theory and Experiment. Working paper
- Szkup, Michal, & Trevino, Isabel. (2020). Sentiments, strategic uncertainty, and information structures in coordination games. *Games and Economic Behavior*, 124, 534–553.
- Weg, Eythan, Zwick, Rami, & Rapoport, Amnon. (1999). An outside option experiment. *Games and Economic Behavior*, 14, 150–286.
- Yang, Ming. (2015). Coordination with flexible information acquisition. *Journal of Economic Theory*, 158, 721–738.
- Yang, Ming. (2020). Optimality of debt under flexible information acquisition. *The Review of Economic Studies*, 87(1), 487–536.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.