On The Cognitive Foundations of Trade*

Jacopo Magnani Ryan Oprea†

May 13, 2026

Abstract

A recent literature has shown that, in individual choice settings, cognitive frictions often cause people to (i) imprecisely process information, (ii) imperfectly allocate attention and (iii) incompletely assess economic decisions. We study the implications of these canonical frictions for markets by evaluating to what extent they explain widely-observed speculative trade in violation of "no-trade" theorems. Diagnostically stripping away these frictions in experimental exchange tasks, we find that all three play a role in inducing trade, but to varying degrees. By contrast, overconfidence (a popular alternative behavioral explanation) does not drive trade in our setting. We show that some frictions are amplified by equilibrium effects in interactive exchange.

Keywords: Cognitive Economics, No-Trade Theorems, Financial Markets, Overconfidence, Behavioral Economics, Experimental Economics

JEL codes: C9, D03, G1

1 Introduction

A recent literature in behavioral economics has identified a collection of *cognitive frictions* that generate systematic behavioral departures from standard predictions (Enke, 2026). This literature

*Ryan Oprea acknowledges financial support from the National Science Foundation under Grant SES-2619167. The experiment was approved by the Human Subjects Committee of the University of California, Santa Barbara. This paper replaces and subsumes our earlier working paper, "Why Do People Violate No Trade Theorems? A Diagnostic Test."

†Magnani: Economics Department, Norwegian University of Science and Technology, Trondheim, Norway, jacopo.magnani@gmail.com; Oprea: Haas School of Business, University of California, Berkeley, Berkeley, CA, 95064, roprea@gmail.com.

has shown that humans tend to (i) imprecisely process information, (ii) imperfectly allocate their attention and (iii) incompletely assess the structure of choices, and that these tendencies can explain many important anomalies in economics. Most of this literature has so far been focused on understanding *individual decision making* and has argued that these cognitive frictions can explain and unify our understanding of key choice anomalies in the domains of risk, time and belief formation. In this paper, we take some early steps in understanding the role these same cognitive frictions play in shaping *market behavior*.

We experimentally examine to what extent cognitive frictions identified in behavioral economics can explain the puzzling human propensity to *trade* even when it is not rational to do so. Some trade, of course, is perfectly rational and requires no explanation: differences in valuations for consumables, for instance, are sufficient to explain why people exchange them. However when the value of the traded object is uncertain and similar across traders – for instance in the exchange of assets in financial markets – the human propensity to trade is difficult to explain. As a series of "No-Trade" theorems show (e.g. Milgrom and Stokey (1982), Tirole (1982)), rational traders should not engage in this kind of speculative trade at all: the mere fact that one's trading partner is *willing* to sell such an object transmits enough negative information on the object's value to make trade irrational. Nonetheless, robust financial markets built on exactly this type of speculative exchange are central to most economies – one of the key facts to be explained in market microstructure. The literature has offered a number of institutional explanations (discussed in more detail below) for this puzzling phenomenon, but robust speculative exchange continues to arise in settings where such explanations are unavailable, including in controlled laboratory studies (Plott and Sunder (1988), Angrisani et al. (2008), Carrillo and Palfrey (2011)).

To study whether (and which) cognitive limitations are responsible for anomalous trade in these settings, we focus on arguably the simplest possible setting of market exchange – one in which the relevant forces shaping trade are cast in particularly sharp relief. Two decision makers are each endowed with a different object (red or blue), each with an unknown (but common) value to the decision makers. Each decision maker is also endowed with a noisy private signal of the value of their own endowment. The two decision makers each simultaneously decide whether to swap their objects with one another, each incurring a small cost if they choose to do so. If both parties agree to the trade, it is consummated. In this simple setting, standard no-trade theorems predict that neither agent (if rational) should be willing to trade. Intuitively, if one's counterpart is willing to trade, it suggests they have negative information on their own item's value, making the trade unappealing.

We investigate the role three canonical cognitive frictions – identified in the recent behavioral economics literature – play in generating trade in this simple setting. We describe these frictions in terms of *strategies* that they naturally inspire – decision procedures that simplify away aspects of the optimal strategy that generate specific cognitive costs or run afoul of specific cognitive constraints. Arguably the three most important such strategies in this literature are the use of:

1. **Imprecise strategies** in which agents noisily process and respond to information, economizing on cognitive costs associated with precise evaluation and decision.

2. **Inattentive strategies** in which agents focus attention on salient information (information that naturally "stands out") at the expense of less salient (but potentially equally important) information, economizing on cognitive costs associated with deliberately directed attention.

3. **Incomplete strategies** in which agents ignore key steps of reasoning or pieces of information in the problem altogether, economizing on cognitive effort costs associated with computation.

By imprecisely, inattentively or incompletely processing information, people can significantly reduce the cognitive burden of decision-making, but often at the expense of decision quality. As we discuss in more detail below, such strategies have been documented in a wide variety of individual decision settings including decision under risk, intertemporal choice and Bayesian inference in the recent literature.

We first show theoretically that variations on *each* of these broad cognitive strategies can generate violations of no-trade theorems, inducing people to trade when they shouldn't. *Imprecise strategies* induce random behavior, turning people effectively into (literal) "noise traders" for purely cognitive reasons. Randomization not only induces trade directly but also (as the microstructure literature has emphasized, e.g. Kyle (1985) and Glosten and Milgrom (1985)) can create motivation for *others* to trade as well via equilibrium effects. *Inattentive* strategies, on the other hand, generate scope for trade because (as recent research has documented, e.g., Conlon et al. (2022)), people often treat their own information as more salient than others' information, causing one's own information to have a stronger influence on behavior. If agents put sufficiently more weight on their own information than others' when forming beliefs, they will be induced to trade in violation of no-trade theorems. Finally, in our simple setting the only step of reasoning that can be removed to simplify decision-making is contingent reasoning about the implications of others' willingness to trade. Because of this, *incomplete strategies* that ignore this step of reasoning (and therefore the information contained in others' willingness to trade) will generate irrational trade. We show that these three frictions generate willingness to trade not only in isolation but also in equilibrium in strategic extensions (which are, themselves, formally equivalent to behavioral game theory models like quantal response and cursed equilibrium models).

To empirically study the role these strategies (and the frictions that motivate them) play in inducing trade, we designed and ran an experiment based on the simple trading game described above. In our baseline exogenous information ("EX") tasks, subjects are assigned abstract objects (red or blue) that potentially have value, are matched into pairs and are each given binary noisy signals on whether their or their counterpart's object is valuable. They then make a binary decision on whether to trade, with trades only being executed if both parties agree to do so. In one task, using the strategy method, we have subjects make this decision contingent on either of the two signals they might receive. In another, we have them make the decision contingent on each possible combination of their *and* their counterpart's signal, removing the need to contingently reason about the decision. We show that the resulting six decisions allow us to diagnostically strip away the availability of the cognitive strategies discussed above, allowing us to gather bounds on the rate at which each of these strategies is used among subject traders.

We find that when all three strategies are available, most subjects (66%) choose to trade, violating no-trade theorems in even the simplest possible markets. By contrast, when scope for inattentive and incomplete strategies are removed and only imprecise strategies remain to motivate trade, only 10-20% choose to do so. Using the diagnostic structure of our design to decompose behavior, we find that all three cognitive strategies contribute to the high rate of trade in this environment, though to different degrees. We estimate that incomplete strategies are most commonly used by subjects, with between 33% and 54% of subjects using incomplete strategies. Imprecise strategies are estimated to be used by 33% of subjects. Finally, inattentive strategies are used by between 14% and 32% of subjects. Many subjects appear to use more than one of these simplification strategies in concert.

We complement this reduced form analysis with a structural analysis that confirms the importance of all three strategies but gives us a more nuanced picture of their effects. Counterfactuals based on structural estimates suggest that levels of both the imprecise and incomplete strategies used by subjects are strong enough to generate significant trade in isolation (i.e., without the complementary use of the other strategies), while the estimated use of inattentive strategies is not strong enough to sustain trade on its own. In particular, estimates suggest that incomplete strategies are strongest in isolation (capable of generating 100% trade rates in subjects), imprecise strategies are somewhat weaker (capable of driving around 50% trade rates) and inattentive strategies weakest (unable to sustain any trade on its own). The estimates also suggest that these strategies are strong substitutes for one another: removing only one strategy when the others are in use has little effect at reducing propensities to trade when the other two are also in use. Given the fact that many subjects appear to be prone to the use of multiple simplification strategies, this suggests that interventions aimed at reducing irrational trade need to be tailored to target multiple frictions at once.

Crucially, our structural estimates also allow us to assess the role strategic *equilibrium effects* play in generating the trading behavior we observe. We find evidence that these equilibrium effects influence trade, above and beyond the direct influence of these cognitive strategies: several lines of evidence suggest that our data is better explained by models that include equilibrium beliefs than by models that include only natural, naive beliefs. Estimates that include scope for equilibrium effects suggest that imprecise strategies have a larger effect on trade and inattentive and incomplete strategies a smaller effect than estimates that do not account for such effects. These results highlight the importance of taking strategic reasoning and equilibrium amplification effects into consideration when assessing the impact of cognitive frictions on behavior. They also emphasize the connection between cognitive frictions recently discussed in the literature and core models from behavioral game theory (like QRE and cursed equilibrium models).

Our experiment allows us to rule out a number of potential alternative explanations for our results, allowing us to better understand the mechanisms driving trade in our data. First, by design, our experiment rules out risk aversion as a driver of trade ex ante. Second, our results are nearly identical when we remove subjects who made mistakes on comprehension questions, ruling out the possibility that our results are driven by confusion. Third, removing subjects who have low opinions of their counterparts' intelligence has little effect on our results, suggesting that our results aren't driven by dismissive beliefs about counterpart rationality (as in, e.g., some alternative interpretations of "cursed reasoning" models). This underscores that our "incomplete strategy" results really are driven by a failure to reason fully rather than by suspicions about the value of the information contained in the counterpart's actions. Finally, we find that rates of trade in our setting are actually *higher* among our most cognitively capable subjects (according to standard cognitive tests), suggesting that use of these strategies may be a deliberate effort to economize on cognitive costs rather than a consequence of binding cognitive constraints.

Finally, in order to examine the robustness of our findings, and benchmark their magnitude relative to other salient potential psychological drivers of trade, we replicate our entire design in a setting in which subjects' own cognitive abilities influence the precision of their signals. In our *endogenous signal* ("END") condition, the quality of traders' information depends on their performance in an IQ test. Not only does this add ecological realism to the experimental task, it creates scope for an alternative potential psychological driver of trade: relative overconfidence (Daniel and Hirshleifer, 2015) . We theoretically show how overconfidence in one's own signals can induce trade and how this can be detected in the data.

We find that trade increases in relevant END tasks in a manner that mirrors the predicted effects of overconfidence. Importantly, however, when we directly elicit subjects' beliefs about their relevant cognitive skills relative to their counterpart (and therefore their beliefs about the relative quality of their signals), we find that there is no relationship between relative overconfidence and this increase in trade. Instead, endogenous signals appear to intensify the use and strength of *inattentive strategies*: when signals are endogenous, one's own signal becomes relatively *more salient*, causing inattentive subjects to overweight them more, generating a more powerful influence of inattentive strategies on trade. This suggests that effects that mirror overconfidence in our data are generated for reasons that have less to do with miscalibrated beliefs than with poorly allocated attention. It also suggests that inattentive strategies may be more important in driving trade in real-world markets (where signals often *are* the endogenous outcome of deliberation) than our EX condition results would suggest.

To summarize, we show that a prominent collection of cognitive frictions, recently emphasized in the behavioral literature, conspire to produce high rates of anomalous trade in speculative exchange. While we know much from the recent literature about how these frictions generate anomalies in individual choice, we know less about their implications for markets. Our results show that these frictions, collectively, generate significant anomalous trading behavior in even very simple markets, producing dramatic failures of standard theoretical predictions. Our findings point to the potential value of recent work in cognitive economics for understanding larger scale social interactions like market exchange. Indeed, we find evidence that the frictions documented in this literature are sometimes amplified by equilibrium effects in interactive trading decisions. Given that naturally occurring markets are far richer and more complex than those studied in our experiment (and the scope for these frictions therefore potentially far more severe), we think our results point to the likely importance of cognitive limitations in helping to explain the active speculative trade we observe in, e.g., real-world financial markets.[1]

[1] Methodologically, our results also suggest a potentially valuable caution in the interpretation of behavioral data. Our results include some behaviors that could easily be attributed to systematically distorted beliefs but are instead driven by broader, domain-general cognitive frictions. In particular, had we not measured subjects' beliefs about the cognitive abilities of their counterparts, we could easily have viewed our evidence on incomplete strategies as instead

The remainder of our paper is organized as follows. After reviewing the related literature in Section 1.1, in Section 2.3 we describe our setting, prove a no-trade theorem and show how some recently discussed cognitive frictions from the literature can generate anomalous trade in this setting. In Section 3, we discuss our experimental design and its diagnostic structure. In Section 4, we discuss our main results from our EX condition and in Section 5 show some additional results that help us better understand the mechanism behind our findings. In Section 6, we discuss results from our END treatment and in Section 7 report our structural estimates. We conclude with a discussion in Section 8.

1.1 Related Literature

Our paper relates to several strands of literature. The first is a growing literature that studies the role a small set of cognitive frictions play in generating a wide range of behavioral anomalies. See Enke (2026) for a recent review. One prominent friction is *noisy or imprecise processing*, which leads to attenuated responses of beliefs and choices to underlying fundamentals. Probability weighting, for example, can be understood as a consequence of such attenuation: individuals respond too weakly to differences in objective probabilities (Enke and Graeber, 2023; Frydman and Jin, 2023; Oprea, 2024; Enke et al., 2024). Similarly, hyperbolic discounting over long horizons can be interpreted as an attenuation effect in the representation of future outcomes (Enke et al., 2025; Ebert and Prelec, 2007). Another class of cognitive frictions operates through *selective attention*, whereby individuals focus on only a subset of available variables. In models of salience, attention is drawn in a bottom-up manner to attributes that are prominent due to contrast or surprise, leading decision makers to distort the relative importance of different features. As discussed in Bordalo et al. (2022), this mechanism helps explain anomalies in choice under risk (Bordalo et al., 2012b), the endowment effect (Bordalo et al., 2012a), and decoy effects and context-dependent willingness to pay (Bordalo et al., 2013). Salience-based distortions can also account for anomalies in judgment and inference, such as over- or underreaction in belief updating (Enke and Graeber, 2023; Augenblick et al., 2025; Ba et al., 2025; Bordalo et al., 2026). A third class of frictions arises when agents effectively solve a simpler, misspecified version of an original decision problem that omits key contingencies, leading to *failures of contingent reasoning* (Esponda and Vespa, 2023; Niederle and

an outgrowth of dismissive beliefs about the ability of counterparts (a'la a literal reading of a cursed equilibrium model). Likewise, had we not measured belief data directly, we could easily have mis-read the increase in attachment to self-generated signals as evidence of over-confidence instead of an intensification of generic salience effects. Our results may therefore point to the value of eliciting beliefs directly in experiments to separate the causal effect of beliefs from those of cognitive frictions.

Vespa, 2023). This mechanism has been used to explain a wide range of phenomena, including the winner's curse in auctions (Charness and Levin, 2009; Martínez-Marquina et al., 2019), mistakes in voting (Esponda and Vespa, 2014), failures to learn from market prices (Ngangoué and Weizsäcker, 2021). In a diagnostic decomposition of the winner's curse, Nagel et al. (2025) show that failures of contingent thinking are of first-order importance relative to computational mistakes and non-standard preferences. Our contribution is to explore how these same cognitive frictions shape market behavior, and in particular speculative efforts to trade on the basis of private information.

Second, our paper contributes to a literature attempting to explain why exchange and price formation occurs in financial markets despite no-trade theorems. While some explanations for trade emphasize rational forces such as portfolio rebalancing and agency motives (e.g. Dow and Gorton (1997)), the volume and nature of trade observed in financial markets are difficult to reconcile solely with these factors.[2] This suggests that behavioral forces are likely to play an important role. A large literature has proposed such behavioral explanations. A first mechanism is *noise*. The "noise trader" approach (Shleifer and Summers, 1990) posits that investors often trade on pseudo-signals that are uncorrelated to fundamentals (such as the advice of financial commentators). In turn, the presence of noise traders can relax adverse selection concerns and thereby create incentives for others to trade on their information (e.g. Kyle, 1985; Glosten and Milgrom, 1985). A second mechanism is *cursed reasoning*. In models of cursed beliefs, agents fail to fully account for the fact that others' willingness to trade is conditioned on their private information, leading to excessive trade in equilibrium (Eyster et al., 2019). A third mechanism is *biased attention*. A growing empirical literature shows that investors' trades are driven by attention-grabbing events. For instance, individual investors are more likely to buy attention-grabbing stocks (Barber and Odean, 2008) and are disproportionately likely to sell both extreme winners and extreme losers in their portfolios (Hartzmark, 2015). A fourth mechanism is *relative overconfidence*. If investors believe their information is more accurate than that of their trading partners, the common priors assumption underlying no-trade theorems is violated, generating scope for trade (e.g. Odean 1998; Daniel et al. 2001; Scheinkman and Xiong 2003; see Daniel and Hirshleifer 2015 for a review). Our contribution is to provide a unified experimental framework that allows us to study these four mechanisms within a common environment. By designing treatments that selectively isolate

[2] For instance, Ross (1989) writes: "It is difficult to imagine that the volume of trade in security markets has very much to do with the modest amount of trading required to accomplish the continual and gradual portfolio balancing inherent in our current intertemporal models" and, noting over 200% market turnover in 2007, French (2008) argues that "[f]rom the perspective of the negative-sum game, it is hard to understand why equity investors pay to turn their aggregate portfolio over more than two times in 2007."

noise, cursed reasoning, limited attention, and overconfidence, we are able to assess their relative importance in generating trade.

Third, our research builds on a small prior experimental literature studying no-trade theorems.[3,4] Two pioneering papers are particularly relevant: Angrisani et al. (2008) and Carrillo and Palfrey (2011) both study bilateral trading interaction with endogenous prices between privately informed buyers and sellers.[5] Both document frequent trade in settings where no-trade theorems suggest it should not occur and both provide some clues as to potential sources of these violations. Carrillo and Palfrey (2011) show that the frequency of violations varies with the type of trading mechanism (auction vs. posted price) and provide some *ex post* evidence that results are broadly consistent with cursed equilibrium.[6] Angrisani et al. (2008) show that long run learning can substantially ease violations of no-trade theorems. Our experiment differs from this previous literature by (i) studying a game that simplifies the analysis of no-trade violations [7] and (ii) by designing our experiment to identify alternative cognitive sources of violations of no-trade theorems. Our results

[3]The earliest experimental evidence suggesting potential violations of no-trade theorems that we are aware of is in the literature on information aggregation in markets. For example, in one treatment of the seminal experiment reported in Plott and Sunder (1988), privately informed subjects trade in a double auction even though the asset yields the same state-contingent dividend to all players.

[4]We previously conducted a related experiment studying trade under private information using a similar basic environment but a different implementation and treatment design. A summary and a comparison with the current experiment is provided in Appendix D.

[5]Other papers have produced evidence of non-equilibrium behavior in different private-information games, such as betting games (Brocas et al. 2014), common value auctions (see in particular the recent literature on the maximal auction: Ivanov et al. 2010, Camerer et al. 2016), common value elections (Esponda and Vespa 2014) and the compromise game (Carrillo and Palfrey 2009).

[6]Other work such as Esponda and Vespa (2014) and Ngangoué and Weizsäcker (2021) have reported evidence that subjects make systematically different choices in simultaneous-move and sequential-move games with private information. Indeed, subjects seem to be better at inferring information from other players' choices when the outcomes of these choices are directly observed rather than left to be hypothetically inferred. This suggests the sort of limits in hypothetical or conditional thinking formalized in cursed equilibrium models.

[7]Specifically, we study a game that is simpler than the ones studied in prior work, with fewer stochastic and endogenous components, making it somewhat easier to interpret results in the light of potential explanations. One major difference in our game is that potential values of assets are binary rather than continuous. Another is that subjects do not form prices but simply make a binary decision to trade (or not). Both of these simplifications are necessary for our diagnostic treatment design, but they have the added benefit that they do not require us to condition our results on realizations of continuous random variables or endogenous price choices by subjects. Instead, we vary the costs and benefits of trade directly and exogenously by varying transaction fees subjects face by trading, giving us somewhat easier to interpret evidence on the motivations behind trade. Finally, because subjects directly trade assets with identical risk characteristics, our design eliminates risk or loss postures as explanations in a very direct and (to subjects) transparent way.

are broadly consistent with the prior literature and help to explain many of the findings in this prior body of work.

2 Conceptual background

In Section 2.1 we describe and model the simple trading environment that we implement in our experiment and prove a no-trade theorem for this environment. In Section 2.2 we discuss three potential cognitive drivers of violations of the no-trade theorem and state a set of supporting Propositions. All proofs are collected in Appendix A.

2.1 Model and No-Trade Theorem

Consider two agents, $i \in \{1, 2\}$, each of whom is endowed with an Arrow security whose value depends on the state $\omega \in \{1, 2\}$. Agent i's security pays $v = \overline{v}$ if the state is i and $v = \underline{v}$ if the state is $j \neq i$, with $\overline{v} > \underline{v}$. Each state occurs with equal probability and each agent i receives a private signal $s_i \in \{1, 2\}$. For simplicity we assume that the two signals are independent conditional on the state. We define $\mu \in (0.5, 1)$ as the likelihood that the signal player i receives is correct:

$$\mu \equiv \Pr(s_i = k|\omega = k) \geq 0.5 \text{ for } i \in \{1, 2\}, k \in \{1, 2\}$$

Each agent i chooses, simultaneously, whether or not to agree to trade Arrow securities with their counterpart by choosing $a_i = T$ (agree to trade) or $a_i = N$ (do not agree to trade). We denote player i's counterpart by $j \in \{1, 2\}$, $j \neq i$. If both players choose T then the two agents exchange assets and each pays a transaction fee f (the fee is a pure cost, not a transfer to the other agent). After making trading decisions, uncertainty about the state is resolved and payoffs are realized. We assume that each player i has a strictly increasing Bernoulli utility function, $u_i(\cdot)$, over these payoffs.

At the core of agent i's trading decision is the expected gain from trading relative to not trading conditional on $s_i = k$, denoted $\Delta_{i,k}$:

$$\Delta_{i,k} \equiv \tau_{i,k} \left\{ [\pi_{i,k} u_i(\overline{v} - f) + (1 - \pi_{i,k}) u_i(\underline{v} - f)] - [\pi_{i,k} u_i(\underline{v}) + (1 - \pi_{i,k}) u_i(\overline{v})] \right\}$$

where $\tau_{i,k} \equiv \Pr(a_j = T|s_i = k)$ is player i's conditional belief that her counterpart will trade (i.e. that $a_j = T$) and $\pi_{i,k} \equiv \Pr(\omega = j|s_i = k, a_j = T)$ is the probability she attaches to the state being j (i.e. that it is advantageous for her to trade) upon receiving signal $s_i = k$ and conditional on her counterpart trading.

Whenever agents face a transaction fee greater than zero, a version of the no-trade theorem applies to this game: there is no Bayes-Nash equilibrium in which players trade. As with other no-trade theorems, the intuition is that, given the common value nature of the assets, neither player can expect to gain from a trade. The intuition is most easily described for the case where each player i plans to trade upon receiving signal $s_i = j$ (but the proof is more general). In this case a trade can only occur if agents have received opposite signals and since these signals are equally accurate in expectation, an agent should believe each security is equally likely to be valuable. Conditional on a trade occurring, therefore, both agents will, in expectation, lose the amount of the transaction fee by trading and therefore should never agree to trade in equilibrium. Importantly for our experiment, this no-trade result holds for any utility functions and is thus robust to risk-aversion and loss-aversion.

Proposition 1. *Whenever $f > 0$, at least one player chooses N for each signal in any Bayes-Nash equilibrium. Thus trade cannot occur in equilibrium.*

2.2 Three Simplification Strategies

Although trading optimally in this minimal environment may seem relatively simple, it nonetheless requires traders to marshall significant cognitive resources. In particular, trading rationally requires decision makers to exert (i) significant cognitive control to make **precise**, internally consistent decisions, (ii) properly allocated **attention** to optimally weight the various pieces of information embedded in the problem and (iii) significant **computational effort** in order to properly extract information on counterparts' signals from their behavior under various contingencies. As literatures in cognitive science and neuroscience emphasize, each of these elements of cognition (precision, attention and computational effort) require the use of constrained resources including executive cognitive control, working memory, representational precision and time. Optimal behavior, in other words, is cognitively *costly.*

A growing literature in behavioral economics and cognitive science has documented that decision makers often respond to these kinds of cognitive costs (i.e., cost of optimal behavior) by substituting to simpler-than-optimal decision strategies that economize on them. We emphasize three especially important simplification strategies that have been the focus of much attention in the recent literature. We emphasize that these strategies each economize on an element of reasoning required to trade rationally, and thus show that instances of each these strategies can generate violations of the no-trade theorem discussed in the previous subsection. Although we will use the word "strategy" throughout to describe these boundedly rational behaviors, we do not claim they are necessarily deliberate – they may simply represent default rules used when the agent is unable or unwilling to expend the resources necessary to identify and implement the optimal choice rule.

Imprecise strategies. First, as a long literature in neuroscience has emphasized, computing, perceiving and representing information in one's brain *precisely* is metabolically expensive (see for example the textbook treatment of Sterling and Laughlin, 2015). People often respond to this expense by **imprecisely** representing information with noise, reducing the costs of responding to the environment, often generating imprecise (i.e., noisy) behavior as a consequence. In perceptual tasks, for instance, a long line of literature documents that decision-makers tend to store or represent visual and auditory information imprecisely; see the review of key psychophysics studies in Woodford (2020). A recent literature in economics and cognitive science shows that this tendency to use imprecise cognitive strategies extends also to decision making: decision makers tend to represent numerical computations imprecisely, generating sometimes systematic distortions in behavior (e.g. Polanía et al., 2019; Khaw et al., 2021; Frydman and Jin, 2022; Enke and Graeber, 2023).

Importantly, noisy cognition often is accompanied by *behavioral* noise (e.g. Khaw et al. (2021); Enke and Graeber (2023)). To the degree this is true, in our simple, binary trading environment, substituting from a precise, optimal strategy to an imprecise strategy will *immediately* generate more trade than the zero trade predicted for precise decision-makers.

Inattentive strategies. Second, another literature in cognitive psychology and economics emphasizes that attention is a constrained (and costly-to-deploy) resource, and that people often economize on cognitive effort (i.e., simplify their task) by limiting the amount of attention they devote to decision-making (Loewenstein and Wojtowicz, 2025). One important effect of using limited attention strategies, is that it causes decision-makers to automate the weight they apply to various pieces of relevant information in the environment. For instance, an important recent literature emphasizes that **inattentive strategies** tend to cause decision-makers to focus their limited attention on *salient* pieces of information – information that "stands out" for reasons that are not always rational (see Bordalo et al., 2012a,b, 2013, 2022, 2026). As a consequence, inattentive strategies often cause decision makers to overweight salient information relative to rational benchmarks.

In our trading environment, making an optimal decision requires decision-makers to equally weight two key pieces of information: their own signal and their counterpart's signal. However, there is significant evidence from previous experiments (e.g. Nöth and Weber, 2003; Goeree et al., 2007; Weizsäcker, 2010; Conlon et al., 2022) that people's own signals tend to be more salient than signals assigned to others. Why do people behave as if their own signals are particularly salient? A recent literature has shown that bottom-up attention can be driven by objective factors (such as the diagnosticity of signals) as well as subjective memories (e.g. Gennaioli and Shleifer, 2010; Bordalo et al., 2023). In our context, since typically individuals do not observe the private information of others, past experiences are likely to highlight the predictive power of one's own information, rather than that of others' private information, for making judgments (e.g. about the profitability of trades). This leads to a focus on one's own private signal.

To the degree this is true in our environment, we would expect the use of inattentive strategies to lead people to apply more focus and consequently more weight to their own signals relative to their counterparts. Such asymmetric overweighting in turn will tend to causes subjects to neglect other traders' information, leading them to trade even when our no-trade theorem suggests they shouldn't.

Incomplete strategies. Finally, making rational economic decision demands significant computational effort, often requiring the decision maker to engage in several steps of reasoning in order to form a rational decision. This computation is costly both in terms of time required and the disutility of the effort involved. Thus, a third way of economizing on cognitive costs is to reduce these computational burdens by eliminating computational steps from the task, ignoring relevant components of the problem altogether and solving a simpler version of the problem instead. By incompletely solving the problem, the decision maker removes whatever costs would have been expended in the eliminated steps. This type of simplification is closely related to the notion of attribute substitution (Frederick and Kahneman, 2002), whereby individuals replace a complex judgment with an easier one, as well as to models of sparse cognition (Gabaix, 2014), in which agents selectively ignore parts of the decision problem.

Arguably only one step in our simple trading problem can be cleanly edited from the decision-making process and ignored altogether to produce an incomplete strategy of this sort. Specifically, in addition to interpreting her own signal, the trader optimally must reason about the signal implied by her counterpart's action and determine under what contingencies those inferences are relevant. Thus, the most natural way to simplify the computations involved in the trade decision is simply to ignore this inference problem altogether, ignoring the possibility that one's counterpart's actions contain any useful information. Doing so removes *most* of the complicated reasoning required in the task, making the computations involved in the incomplete strategy that remains trivial. This kind of incomplete reasoning has been found in other complex environments that require thinking about unrealized contingencies, such as choices under risk and uncertainty, elections and auctions (see the review of Niederle and Vespa, 2023). In financial markets as well as our experiment, incomplete strategies will induce people to trade in violation of no-trade theorems.

2.3 Theory: Simplification Strategies and Trade

We formally model the tendency of these three simplification strategies to induce trade, in two distinct ways (details and proofs are in the Appendix). First (and almost trivially) we verify that each of these strategies will tend to immediately produce trade in violation of the no-trade theorem. Second, we also show that they will tend to produce trade in equilibrium in standard behavioral game theory models. In some case, this will occur due to *equilibrium effects*: knowing that one's counterpart might trade against her information eases the adverse selection in the problem, inducing trade in otherwise rational players. We will use these equilibrium models (and contrast them to non-equilibrium variations) in our structural estimation in Section 7.

To do this we assume players are risk-neutral. Under risk-neutrality, we can simplify the expression for player i's expected gain from trading as follows:

$$\Delta_{i,k} = \{\tau_{i,k} \left[2(\overline{v} - \underline{v})(\pi_{i,k} - 0.5) - f\right]\} \tag{1}$$

We emphasize, however, that all of the arguments discussed here hold for any monotonic Bernoulli utility function.

We begin the analysis by treating player i's beliefs about the strategy of his counterpart as exogenous, and require them to be determined in equilibrium only in a second step. We denote player i's belief that player j will trade conditional on signal $s_j = l$ by $b_{j,l}$. Starting from these beliefs, a rational player will form a posterior belief about the state conditional on the observed private signal $s_i = k$. We focus on the contingency $s_i = j$, since that is the case where deviations from rational behavior matter most. We denote the rational posterior by $\pi^*_{i,j}$.[8] As shown below, some of the simplification strategies can drive a wedge between the actual posterior $\pi_{i,j}$ and the rational posterior $\pi^*_{i,j}$. We restrict the analysis to the case where player i believes the strategy of player j induces sufficient adverse selection so that he would find it unprofitable to trade at the current transaction fee level in the absence of cognitive frictions: beliefs $b_{j,l}$ are such that $\Delta_{i,j} < 0$, that is $\pi^*_{i,j} < 0.5 + \frac{f}{2(\overline{v} - \underline{v})}$.

[8]The rational posterior is $\pi^*_{i,j} = \frac{\mu\left[b_{j,i}(1-\mu)+b_{j,j}\mu\right]}{\mu\left[b_{j,i}(1-\mu)+b_{j,j}\mu\right]+(1-\mu)\left[b_{j,i}\mu+b_{j,j}(1-\mu)\right]}$

Incomplete strategy. First, we formalize incomplete strategies. When player i uses an incomplete strategy to form his posterior belief $\pi_{i,j}$, he partially neglects the contingency in which player j accepts to trade. We assume player i's posterior belief is given by a weighted average of the simple posterior obtained by conditioning only on i's own signal (which simplifies to μ) and the fully rational belief $\pi^*_{i,j}$ which conditions also on the contingency in which a trade occurs:

$$\pi_{i,j} = \alpha\mu + (1-\alpha)\pi^*_{i,j}$$

where $\alpha \in [0,1]$ measures the degree of incompleteness of the mental strategy or the degree of failure in contingent reasoning. When player i ignores counterpart j's action, player i perceives a larger gain from trade than a rational player who conditions on completing the transaction. If the perceived gain is larger than the trading fee then player i expects to profit from a trade. Thus, if the strategy of player i is sufficiently incomplete (α large enough), then trading upon receiving signal $s_i = j$ becomes an optimal strategy.

To extend this idea to an equilibrium setting we further require that player i's beliefs about the strategy of player j are consistent with the actual play of player j. The resulting model is formally equivalent to the cursed equilibrium of Eyster and Rabin (2005), although here we interpret it as a reduced form of failures of contingent reasoning rather than non-equilibrium beliefs (an interpretation Eyster and Rabin (2005) emphasize). Even in equilibrium, sufficiently incomplete strategies can drive trade.

We summarize these arguments in the following Proposition:

Proposition 2.

a) Assume $\mu \geq 0.5 + \frac{f}{2(\overline{v}-\underline{v})}$. *If* $\alpha \geq \frac{\frac{f}{2(\overline{v}-\underline{v})}+0.5-\pi^*_{i,j}}{(\mu-\pi^*_{i,j})}$ *choosing* $a_i = T$ *upon receiving signal* $s_i = j$ *is an optimal action for each player i, given player i's beliefs on player j's strategy.*

b) If $\alpha \geq \frac{f}{2(\overline{v}-\underline{v})(\mu-0.5)}$ *there is an* α*-cursed equilibrium where each player i chooses* $a_i = T$ *if and only if* $s_i = j$.

For example, in the baseline parametrization of our experiment, trading is an equilibrium if $\alpha \geq 0.625$.

Inattentive strategy. Second, we formalize inattentive strategies by assuming each player i overweights his own signal when updating beliefs, since a player's own signal is more salient or top of mind than the counterparts' signal. The posterior belief of player i conditional on signals s_i and s_j is:

$$Posterior_i(\omega|s_i, s_j) \propto \Pr(\omega) \times \Pr(s_i|\omega)^\beta \times \Pr(s_j|\omega)$$

where $\beta \geq 1$ measures the degree of attention-based overweighting. If agents sufficiently overweight their signals, they overreact to their private information when it suggests trading is profitable ($s_i = j$). A similar argument also works in equilibrium (assuming overweighting is common knowledge to close the model).

Formally, if we let $SNR \equiv \frac{\mu}{1-\mu}$, we have the following Proposition:

Proposition 3.

a) If $\beta \geq 1 + \frac{\log\left(\frac{\overline{v}-\underline{v}+f}{\overline{v}-\underline{v}-f}\frac{1+b_{j,j}/b_{j,i}1/SNR}{1+b_{j,j}/b_{j,i}SNR}\right)}{\log SNR}$ *choosing* $a_i = T$ *upon receiving signal* $s_i = j$ *is an optimal action for each player* i*, given player* i*'s beliefs on player* j*'s strategy.*

b) If $\beta \geq 1 + \frac{\log\frac{\overline{v}-\underline{v}+f}{\overline{v}-\underline{v}-f}}{\log SNR}$ *there is a Bayes-Nash equilibrium where each player* i *chooses* $a_i = T$ *if and only if* $s_i = j$*.*

In the baseline version of our experiment, trade can be sustained in equilibrium under inattentive strategies if $\beta \geq 1.6$.

Imprecise strategy. Finally, we formalize imprecise strategies by assuming that player i chooses to trade with positive probability even when the expected gain from trading is negative because of cognitive uncertainty about the actual payoff consequences. Following models of noisy cognition and perception like Woodford (2014) and Matějka and McKay (2015), we model the probability of a trade decision conditional on signal $s_i = k$ as a logistic function of the expected gain from trading:

$$\sigma_{i,k} \equiv \Pr(a_i = T \mid s_i = k) = \frac{e^{\Delta_{i,k}/\gamma}}{1 + e^{\Delta_{i,k}/\gamma}}. \tag{2}$$

Here $\gamma > 0$ is a measure of noise, i.e. attenuation in the responsiveness of a player's choice to the size of the expected gains from trading.[9] When γ is close to zero, behavior is fully responsive to expected payoffs, and since the expected gain from trading is negative (by assumption on player

[9] As suggested by recent studies on noisy perception and cognition, the degree of noise is likely to depend on the characteristics of the choice environment such as time pressure (Polanía et al., 2019), the range of payoffs (Frydman and Jin, 2022) and the complexity of alternatives (Enke and Graeber, 2023). Since we keep these structural characteristics fixed throughout our experiment, we treat noise as an exogenous parameter.

i's beliefs), the probability of trading is also close to zero. For fixed beliefs, as strategies become more imprecise, and γ increases, so does the probability of trading. In the limit $\gamma \to \infty$ behavior is completely unresponsive to expected gains, so that the probability of trading is 50%.

When we apply equilibrium logic by requiring beliefs to be consistent with actual strategies, i.e. $b_{j,l} = \sigma_{j,l}$, this model of imprecise strategies becomes formally equivalent to the logit specification of the quantal response equilibrium (QRE) model (McKelvey and Palfrey, 1995; Goeree et al., 2016), and can generate non-trivial equilibrium effects. Knowing that one's counterpart will agree to trade with some probability regardless of her signal increases the conditional probability that a trade is profitable given that one's counterpart agrees to trade, $\pi_{i,k}$. This, in turn, encourages more trade from i, generating an amplified equilibrium effect. Thus, as strategies become more imprecise, trading probabilities cannot remain arbitrarily small in equilibrium: a sufficiently high γ guarantees that the probability of trading is above a lower bound.

We summarize these arguments in the following Proposition:

Proposition 4.

a) The probability of choosing $a_i = T$ upon receiving signal $s_i = j$ is an increasing function of noise, keeping fixed player i's beliefs on player j's strategy: $\frac{\partial \sigma_{i,j}}{\partial \gamma} > 0$

b) For any $K \in (0, 1/2)$, define

$$\bar{\gamma}(K) = \frac{\overline{v} - \underline{v} + f}{\log\left(\frac{1-K}{K}\right)}.$$

Then, in any quantal response equilibrium with parameter $\gamma \geq \bar{\gamma}(K)$, each player i chooses $a_i = T$ upon receiving signal $s_i = j$ with probability at least K.

This proposition guarantees that, for example, if $\gamma \geq 50$ then in the baseline version of our experiment each player trades upon receiving their signal at least 45% of the time.

3 Design and Empirical Strategy

The experiment consists of twelve decisions, six of which constitute our *exogenous signals* ("EX") condition. We focus attention on the six tasks of the EX condition here, and discuss the other six decisions (which allow us to test for robustness and provide some useful benchmark data) in more detail in Section 6 below.

(a) Screenshot for private signals treatments. (b) Screenshot for public signals treatments.

Figure 1: Strategy choice screenshots.

At the beginning of each game, the computer randomly draws a state $\omega \in \{blue, red\}$ (each occurring with equal likelihood) and assigns each subject an Arrow security that is either of a "blue" or "red" type. A security's value to subjects is $\overline{v}$ =\$10 if it is of the same color as the realized state, and $\underline{v}$ =\$2 otherwise. We do not inform subjects of the state directly but instead give them each a noisy signal which can be either of a "blue" or "red" type as described below. Subjects are each randomly matched with another player, who is endowed with a security of the *opposite color* and given an opportunity to exchange securities with this player. If both subjects in a pair agree to trade, they will exchange securities and they will each pay a transaction fee f =\$2 to the experimenter.

The games elicit choices via the strategy method. After being told their item color (as well as the possible asset payoffs, $\overline{v}$, $\underline{v}$, and the transaction fee, f), subjects make *signal-contingent* choices in each game. The EX condition is divided into two games, visualized in the left and right panels of Figure 1, each an independent version of the game described in Section 2.3. In the *private information game* (pictured in Figure 1a) subjects are simply asked to decide whether they would like to trade as a function of each of the possible noisy signals (called a "clue"), red or blue, they might receive about the higher value security. Thus, subject i chooses $a_i(s_i)$, a mapping $a_i : \{blue, red\} \mapsto \{T, N\}$ from a set of (future) signals to a set of trading decisions (trade, T, or no trade, N). In the *public information game* (pictured in Figure 1b), the subject is instead shown every possible combination of her own and her counterpart's signal, and in each case is asked to make a separate signal-contingent decision. Formally, in these treatments, subject i chooses $a_i(s_i, s_j)$, a mapping $a_i : \{blue, red\} \times \{blue, red\} \mapsto \{T, N\}$. Together then, subjects make six contingent trading decisions in the EX condition (two in the private signals treatment and four in the public signals treatment).

After the experiment is over, the computer implements a trade (and levies a transaction fee on each party) only if both subject i and her counterpart j choose decision functions and receive signals generating outcomes $a_i(s_i) = T$ and $a_j(s_j) = T$. Subjects are given no feedback on the outcomes of trades nor the realizations of states and signals during the experiment.

3.1 Six Diagnostic Decisions

Throughout the paper, we categorize each of the six decisions from the EX condition in terms of how signals should influence the appeal of trade for the subjects. We also assign them labels that specify the simplifications strategy that specifically produces trade in that task, giving a sense of the diagnostic structure of the design.

In *private information* tasks, we will describe the task in which the subject's signal matches her own asset (indication that her asset is valuable and that she should not trade) as $x^?_+$. In this notation, the subscript represents the subject's own signal, while the superscript represent's her counterpart's signal. Thus subscript "+" indicates that the subject's information indicates her own asset is high value; the superscript "?" indicates that the subject *does not know her counterpart's signal*. Likewise $x^?_-$ represents the decision in which the subject receives a signal that does not match her own asset, suggesting that her own asset is the less valuable one which (absent any other inference) indicates that trade is warranted.

Second, in *public information* tasks, we use the same notation but replace "?" with the counterpart's signal explicitly revealed to the subject i. Thus x^-_- is the case in which both subject i's signal and her counterpart j's' signal suggest that subject i would be better off trading and x^+_+ the case in which neither signal suggests this. x^+_- and x^-_+ are the cases in which subject i's and j's signals have opposite implications for i's trade decision.

Thus, our design consists of tasks $x^?_+$ and $x^?_-$ (private information game), and x^+_+, x^-_-, x^-_+, x^+_- (public information game). This set of tasks systematically varies the cognitive strategies available to simplify the trading decision and therefore turns "on" and "off" the availability of the three hypothesized cognitive drivers of violations of no-trade theorems discussed in Section 2.3, above. This allows us to measure the usage, relative importance and effect size of each of these three strategies on trade. To understand the diagnostic nature of the design, we discuss what simplification strategies can drive trade in each of these tasks.

ANY Task. In task $x^?_-$ subjects should not trade, despite receiving information that her asset is the less valuable one. However, the subject has all three simplification strategies available to her and the use of any of them will tend to increase the likelihood of trade:

1. Trivially (as in all of our decisions) she has the option to implement an *imprecise strategy*, injecting noise into the decision which mechanically increases the likelihood of trade relative to no-trade predictions.

2. Because the subject does not know her counterpart's signal, she must infer it, adding a step of reasoning to an otherwise trivially simple task. The only real computational step to remove from the task to simplify it is this inference which will have the effect of encouraging the subject to rely on her signal alone, inducing trade. Thus the use of an *incomplete strategy* induces trade.

3. Finally, even if the subject does use a complete strategy and properly infer her counterpart's signal (conditional on the completion of trade), the subject may use an *inattentive strategy* that fails to properly balance the weight she places on the two signals. Because the subject's own signal is likely to be more *salient*, an inattentive strategy will tend to cause her to overweight evidence that trade is warranted, inducing trade.

Because all three strategies are available here to encourage trade, we will refer to the $x^?_-$ decision as the ANY task.

NOISE Tasks. At the opposite extreme, tasks $x^?_+$, x^+_+ and x^-_+ are ones in which neither incomplete strategies nor inattentive strategies can induce trade for the simple reason that the subject's own private information directly indicates she should not trade. Overweighting this signal relative to her counterpart's will have the effect of reducing the likelihood of trade, meaning inattentive strategies can't induce trade. Likewise, ignoring the counterpart's information will not induce further trade since the subject's own signal already instructs her to do the same thing.

Because trading in these tasks involves trading *against one's own information*, it is difficult to square this type of behavior with anything other than noisy evaluation or behavior.[10] We therefore use these tasks as our most direct piece of evidence of the use of *imprecise simplification strategies*. We will call these our NOISE tasks, and sometimes label $x^?_+$, x^+_+, x^-_+ NOISE?, NOISE+ and NOISE- respectively (acknowledging the subject's information about her counterpart's signal).

[10] One alternative possibility is *confusion* about the primitives of the task. However we are able to rule this out as the primary driver of trade in these tasks (See Section 5 below).

SALIENCE Task. In task x^+_- there is no potential for *incomplete strategies* to induce trade. Like all of our public signals tasks, subjects in this task no longer have to reason towards her counterpart's signal. Instead, the signal is given directly to the subject, removing these costs. However, in this task subjects *do* have to weight her own and her counterpart's signals and trade them off against one another, meaning there is scope for subjects to use *inattentive strategies* that put excess weight on the subject's own, more salient signal. Because this signal indicates low value for her item, it has the potential to produce trade. Because this task includes scope for such salience effects (in addition to noise effects, of course), we call this the SALIENCE task.

INDIFF Task. Least diagnostically useful for our purposes (but included in the design for completeness) is task x^-_- in which it is weakly rational for the subject to trade. We will call this the INDIFF (indifference) task. It is of limited use for us because (i) it is rational to trade in this task but (ii) it is also rational not to trade: subjects shouldn't expect her counterpart to be willing to trade in this condition, meaning subjects have no real incentives to make *either decision*.

3.2 Identifying Effects of Simplification Strategies

The diagnostic idea behind the design is to use variation in rates of trade across these tasks to measure the influence of each of these simplification strategies. In the reduced form analysis, we can do so in the following way:

Imprecise strategies. Trade in any of our three NOISE tasks – NOISE?, NOISE+ or NOISE- ($x^?_+$, x^+_+ and x^-_+) – provide a relatively direct measure of the imprecision in subjects' choices because they each require the subject to trade against the advice of their own signals. Because of the simplicity of our task, random error due to behavioral imprecision seems the most plausible reason for trade in these settings (in Section 5 below, we can also test the alternative possibility that trade in these settings are driven by confusion about the rules of the game).

Inattentive strategies. Trade in our SALIENCE task (x^+_-) could be induced either by inattentive strategies or imprecise strategies. We therefore can measure the amount of trade driven by inattentive strategies by examining the difference in the rate of trade between x^+_- (which include both inattentive and imprecise strategies) and x^-_+ (which include only imprecise strategies), giving us a measure of the amount of trade induced by salient attachment to one's own signal (at least assuming there are no interaction effects between inattentive and imprecise strategies). Another way of seeing this is that x^+_- and x^-_+ are simply permutations of one another that *contain exactly the same information relevant to trade.* If the latter produces more trade than the former, it must be due to an overweighting of one's own salient signal relative to one's counterpart's. Importantly, this is a *conservative* measure of the use of inattentive strategies, since noisy trade in x^-_+ will censor measurement. The difference between x^+_- and x^-_+ gives a lower bound while the rate of trade in x^+_- (SALIENCE) itself gives the upper bound rate of use of inattentive strategies in the data.

Incomplete Strategies. We can measure the influence of incomplete simplification strategies by comparing the rate of trade in the ANY task in which all three strategies are available to the rate of trade in the SALIENCE task in which only inattentive and imprecise strategies are available, netting out the influence of incomplete strategies. Another way of seeing this is that SALIENCE (x^+_-) explicitly provides subjects the exact contingency (the exact pair of signals) that would have to realize in order for a trade to occur in the ANY task ($x^?_-$). It therefore removes the need to perform this computation from the subject, acting as a perfect control for the ANY task that removes precisely the costs of this computation. The difference between these tasks therefore measures the influence these costs have on inducing irrational trade. Again, this is a *conservative* measure, since inattentive trade in SALIENCE (x^+_-) will censor measurement of incomplete strategies. The difference between trade in ANY and SALIENCE ($x^?_-$ and x^+_-) gives a lower bound, while the difference between ANY and NOISE- (x^-_+) gives the upper bound rate of use of incomplete strategies in the data.

Finally, it is worth emphasizing that our model and therefore our experimental design deliberately rules out *risk preferences* as an explanation for trade in all of our treatments. Unlike standard trading games, where players can remove uncertainty and the risk of incurring losses by substituting cash for a risky asset, in our game players can only substitute one risky asset for another. Thus, risk and loss aversion are removed as potential explanations by the design in all of our treatments, making it a particularly sharp instrument for measuring the impact of simplification strategies.

3.3 Robustness and Benchmarking: Overconfidence

In an additional six tasks, we repeat the *EX design* but change the nature of the signals so that their accuracy is linked to subjects' own (independently measured) cognitive ability. We do this for two reasons. The first is simply that it allows us to examine the robustness of many of our main findings to an ecologically realistic setting in which the quality of trade-relevant information is dependent on the trader's ability to process that information. The second, is that it allows us to benchmark the strength of the effects of our three simplification strategy against the effects of a salient alternative mechanism often discussed in the literature: relatively overconfident beliefs (see Daniel and Hirshleifer (2015) for a review of the literature on overconfidence and trade). As we discuss in more depth in Section 6, relative overconfidence in one's own ability to interpret information stands as a belief-based alternative to our cognitive channels for inducing trade.

The six tasks of our *endogenous signals* ("END") condition are identical to those from the public and private information games in our EX condition (indeed, the screenshots from Figure 1 are the same) except for the signal accuracy of the subject's signal, and her counterparts. In particular all subjects complete six Raven's Progressive Matrices (ranging from relatively easy to relatively difficult matrices) and obtain a score $r_i \in \{0, 1, 2, 3, 4, 5, 6\}$. Given subject i's score and the realized state, subject i's signal is drawn from a Bernoulli distribution, $s_i \in \{blue, red\}$, with accuracy $\mu_i = 0.5 + 0.07 \times r_i$. Thus, a subject who obtains a score of 0 is assigned signals with accuracy 0.5, i.e. uninformative signals that are uncorrelated with the state. A higher Raven score results in a higher accuracy. The highest possible accuracy is 0.92. The signals of the two matched subjects in each pair are independent draws with potentially different accuracies, but conditional on the same state. We thus give subjects noisy signals, but signals whose accuracy may vary from subject to subject based on their skill, creating scope for overconfidence.[11] Importantly, the Raven's test takes place *after* subjects submit their strategies, so that subjects are given no feedback on their accuracies. Indeed, as we show below, subjects do not make decisions that suggest that they are aware of their actual accuracy relative to other subjects.

The endogenous signals design gives us (paralleling the notation from the EX condition) tasks $c_+^?$ and $c_-^?$ (private information game), and c_+^+, c_-^-, c_+^-, c_-^+ (public information game). In Section 6 we discuss our diagnostic use of this design in more detail and show that if subjects are overconfident in these tasks, it can create an independent basis for trade.

3.4 Belief Elicitation

Before playing the trading game, subjects completed two belief elicitation tasks regarding performance in the Raven's Progressive Matrices test (the actual test took place at the end of the experiment). Prior to making their forecasts, subjects were shown a brief description of the test

[11] Previous experiments suggest individuals are prone to form overconfident beliefs about their performance in the Raven task, see for example Burks et al. (2013) and Oprea and Yuksel (2022). Kogan (2009) induces skill-based signals in a similar way.

and an example puzzle to familiarize them with the task. The test was described as consisting of six puzzles of varying difficulty, each requiring subjects to select the image that best completes a geometric pattern. Subjects were then asked to forecast how many of the six puzzles they expected to solve correctly, and how many puzzles a randomly selected other participant would solve correctly. Forecasts were elicited as point predictions by selecting an integer between 0 and 6 from a dropdown menu.

The elicitation was incentivized using a straightforward rule: when one of the two belief elicitation tasks was selected for payment, subjects received a bonus payment if and only if their forecast exactly matched the realized score. Simple and transparent incentive schemes of this kind are commonly used in the belief elicitation literature and have been shown to perform well relative to more complex mechanisms (see, e.g., Danz et al., 2022). These belief measures allow us to assess expectations about own and others' ability, which we use to study the role of relative overconfidence in the endogenous signals condition.

3.5 Implementation

We ran the experiment on 250 Prolific participants. All observations were collected on May 4, 2023. The experimental software was coded in oTree (Chen et al., 2016). The experiment pages are reproduced in Online Appendix C. The experiment began with the belief elicitation task. Then subjects were given instructions about the four different trading game treatments and administered a quiz to test and reinforce comprehension. After that, subjects submitted their trading strategies in the four treatments. The order of the four treatments was randomized for each subject. No feedback was provided about the outcome of the games. Finally, subjects answered the six-item Raven's test.

Subjects took on average 13 minutes to complete the experiment. After the experiment was over, subjects were randomly matched in pairs for each game, the skill-based accuracy of each subject was determined (for use in the confidence design tasks), the state and signals realizations were drawn and the outcome in each game determined. All participants received a base payment of around $6, consistent with Prolific policy. One fifth of the subjects were randomly selected to receive a bonus payment that depended on their choices in the experiment. To determine the bonus payment to each eligible subject, one task was randomly selected among the belief elicitation questions, the four games and each of the six Raven's test items (see Azrieli et al. 2018 for a useful discussion on this type of payment protocol). If one of the games was selected, the bonus payment was equal to the subject's payoff realized in that game (thus ranging from 0 to \$10). If one of the Raven's test items was selected, the bonus payment was \$4 if the subject answered correctly and 0 otherwise. If one of the two belief elicitations task was selected, the bonus payment was \$4 if the forecast was correct and 0 otherwise.

4 Results

The left hand panel of Figure 2 plots the raw share of subjects in our *EX condition* who choose to trade in each of the six tasks. The right hand panel plots reduced form estimates of the fraction of trade attributable to each of our three simplification strategies, and the fraction of subjects who use a rational strategy (i.e., who never trade when they shouldn't).

In our baseline ANY task, all three of our cognitive simplification strategies are available to induce trade: (i) subjects receive personal information suggesting their asset is low value (allowing inattentive attachment to salient information to induce trade), (ii) subjects are required to engage in contingent reasoning in order to infer their counterpart's information (allowing incomplete assessment of information to induce trade) and (iii) (trivially) subjects can make noisy decisions that produce trade (allowing imprecise simplifications to induce trade). In this condition 66% of subjects choose to trade in violation of no-trade theorems.

Result 1. *When all three simplification strategies are available to induce trade, 2/3 of subjects choose to trade in violation of no-trade theorems.*

Next, we examine rates of trade when we remove scope for both inattentive and incomplete strategies to produce trade: our NOISE, NOISE+ and NOISE- tasks. In these tasks subjects receive positive signals about their own asset, meaning neither overweighting of own information (inattentive simplification) nor ignoring the information implied by others' actions (incomplete simplification) can generate trade. As the left panel of the Figure suggests rates of trade in these tasks drop to 10-20%, suggesting that behavioral noise is an important component of trade but cannot easily rationalize the nearly 2/3 of subjects who trade in our baseline $x_{-}^{?}$ task.

As discussed in Section 3 we identify a subject as using an imprecise strategy if they choose to trade in any of these tasks. By this measure, as the right hand panel of Figure 2 shows, 1/3 of our subjects can be identified as using inattentive strategies. Most of these subjects only trade in one of the three NOISE tasks, a much smaller proportion in two of them and only 5 subjects in all three, reinforcing our interpretation that trade in these tasks is driven by random errors rather than a

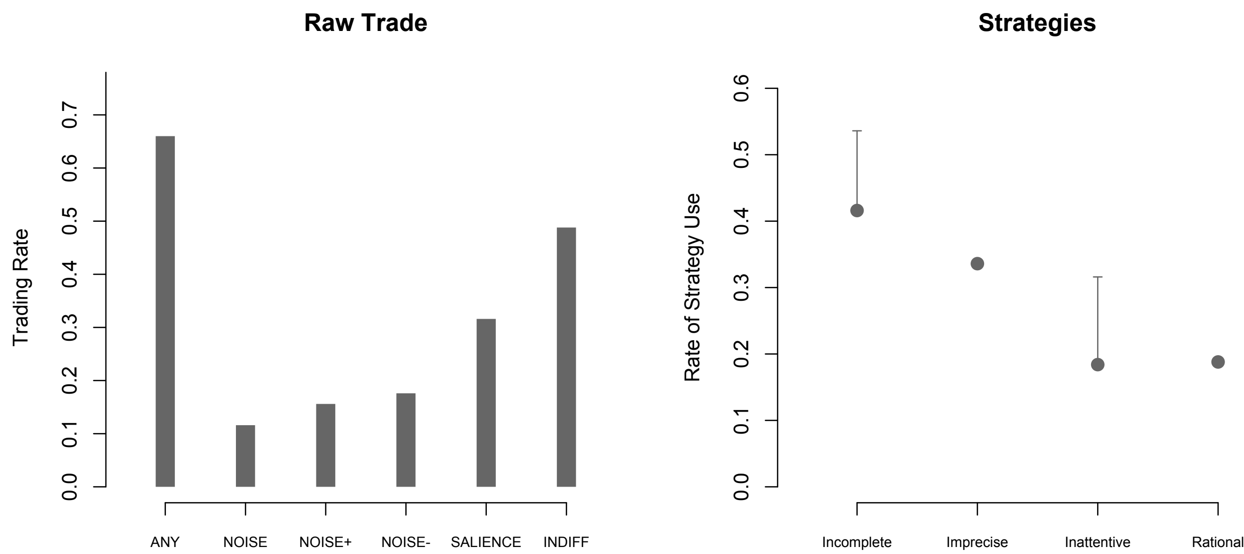

Figure 2: Rates of trade in each task (left) and estimated rates of strategy usage (right) in the full dataset. Vertical bars in the right-hand panel show the upper bound of use of the corresponding strategies.

thorough-going misconception of the task (we provide additional evidence to this effect below).

Noisy behavior can serve as a motivation for trade by others, increase the willingness to trade of other subjects who anticipate noise in the decisions of their trading partners, by reducing adverse-selection in private signal game tasks. Empirically, even when we account for such a channel, the expected gain from trading net of transaction fees when receiving a negative signal remains negative.[12] Thus, the observed rate of imprecise behavior does not seem to be large enough to induce the predominant trade observed in our baseline ANY task. We measure equilibrium effects like these in more detail in our structural analysis in Section 7 below.

Result 2. *33% of subjects trade in at least one of the NOISE tasks, consistent with the use of imprecise strategies.*

Next, we consider to what degree inattentive strategies involving overweighting of salient information in the choice problem generates trade. To do this we compare public information tasks SALIENCE (x^{+}_{-}) and NOISE- (x^{-}_{+}), which are permutations of one another and therefore should produce identical trade for subjects who do not overweight their own information. However, we find

[12]See calculations in Appendix B

that subjects' more salient *own* information has a substantially larger impact on trading behavior than their counterpart's information. Subjects trade 32% in SALIENCE but only 18% in NOISE-. This 14 percentage point difference is statistically significant (with a Wilcoxon test $p < 0.01$). This is a conservative estimate: because noisy trade in NOISE- will censor this difference, this is a lower bound estimate of the use of incomplete strategies. In the right hand panel, a line extending from the estimate shows that up to 31.6% of subjects may use incomplete strategies.

Result 3. *In informationally identical tasks, nearly twice as many subjects trade when their own information supports doing so than when their counterpart's information does. This 14 percentage point difference suggests that inattentive strategies have a significant effect on willingness to trade.*

Finally, we consider how trade changes when we add scope for incomplete strategies that simplify the reasoning problem by ignoring information implicit in the contingency that a trade actually occurs. To measure this, we compare ANY ($x^?_-$) with SALIENCE (x^+_-), which replicates the contingency in which trade occurs in ANY, but which frees the subject from having to reason her way to this information. Trade drops to 31.6% in this latter case, meaning twice as many people trade when forced to engage in computationally intensive contingent reasoning. The difference in the trade frequencies between ANY and SALIENCE, visualized as our conservative estimate of the use of incomplete strategies in the right panel of the Figure, is highly significant (Wilcoxon test $p < 0.01$), suggesting a large effect of incomplete strategies on willingness to trade. Once again, the use of inattentive strategies may censor or measurement of incomplete strategies. The line extending upward from the point shows that the data is consistent with as many as 53.6% of subjects using incomplete strategies (see Section 3 for details).

Result 4. *Twice as many subjects choose to trade on their signal when forced to reason towards their counterpart's information than when this information is transparently provided. This 34 percentage point difference suggests that incomplete strategies have a large effect on the willingness to trade.*

To summarize our reduced form findings, the vast majority of subject trade in violation of no-trade theorems. Diagnostically stripping away the availability of these strategies, we find that incomplete strategies are largest, inducing between 33% and 54% of subjects to trade while imprecise and inattentive strategies each have substantial but somewhat smaller effects.

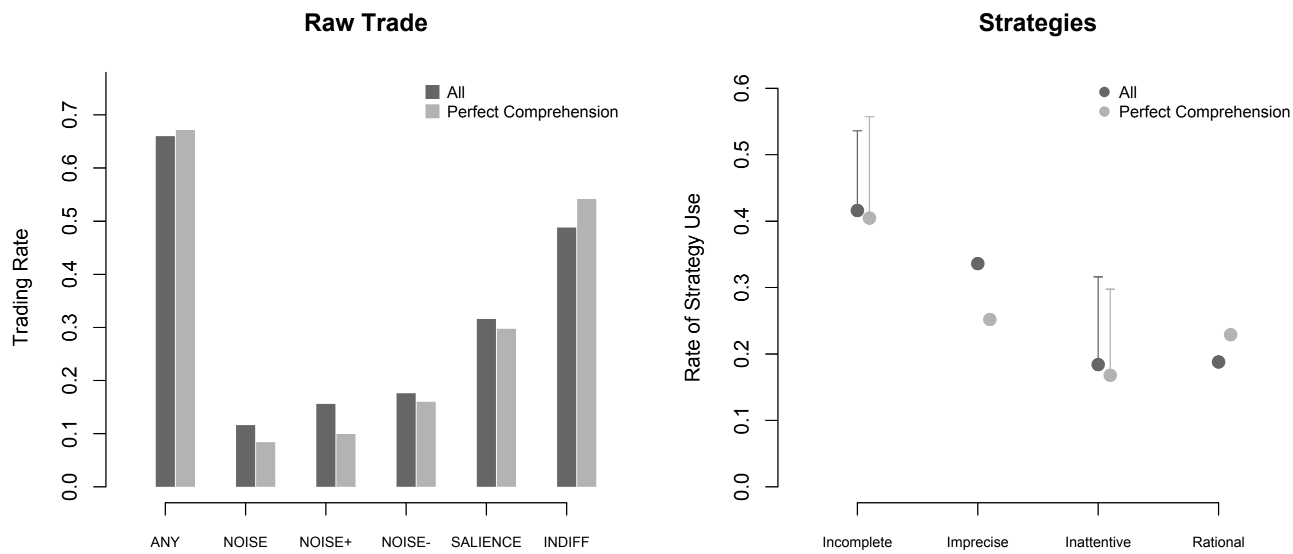

Figure 3: Rates of trade in each task (left) and estimated rates of strategy usage (right) for the full sample (dark gray) and for the subset of subjects who make no errors in the comprehension questions (light gray). Vertical bars in the right-hand panel show the upper bound of use of the corresponding strategies.

5 Evidence on Mechanism

Our motivating hypothesis is that trade is induced in each of these cases by subjects' use of a collection of simplification strategies that economize on cognitive costs and constraints. In this section we provide evidence supporting this interpretation by considering some salient alternatives.

A first alternative possibility is that these findings are not a result of the use of simplified strategies, but instead are driven by systematic confusions about the economic environment induced in the experiment. We attempted to minimize this possibility, ex ante, by studying perhaps the simplest possible economic setting in which no-trade theorems apply. However, we can test this possibility directly by conditioning our assessment on subjects' ability to answer comprehension questions about the rules of the experiment.[13]

[13]This is an imperfect measure of confusion since errors in these questions are surely heavily influenced by inattentive, imprecise answering strategies and incomplete reasoning due to low effort – and since we should expect subjects to learn from feedback on incorrect answers. Nonetheless, if subjects are confused in systematic ways, we would expect it to show up in answers to these questions. In removing subjects who failed to perfectly answer these questions, we are likely removing many inattentive or imprecise subjects (and subjects who learned away initial confusion) but

Figure 3 repeats the data from Figure 2 but adds in lighter gray the same measurements and calculations for the subset of subjects who *perfectly* answer every comprehension question the first time and therefore are very unlikely to be confused about the economic setting of the experiment. As is clear in the Figure, results are very similar in this high comprehension subset, including a slight *increase* in trade in our baseline ALL task. The main difference we observe is a modest reduction in the rate of usage of imprecise strategies (dropping from 1/3 to 1/4 of subjects). But restricting to perfect comprehension subjects has essentially no impact on the rate of use of incomplete and inattentive strategies.

Result 5. *There is little evidence that these results are driven by subject confusion. Subjects who perfectly answered instructions comprehension questions trade at a similar rate, use incomplete and inattentive strategies at nearly the same rate and use imprecise strategies at only a modestly reduced rate.*

A second alternative possibility is that the inattentive and incomplete strategies identified using our ANY and SALIENCE tasks, might be expressions of the imprecise strategies measured independently in our NOISE tasks. In particular, noisy behavior might produce spurious evidence of inattentive or incomplete strategies.

To examine this possibility, we restrict attention only to the subjects who are *not* identified as using imprecise strategies (i.e., who never trade in the three NOISE tasks). The results are plotted in Figure 4, alongside the overall data plotted earlier in Figure 2. Mechanically, rates of trade in the three NOISE tasks in the left hand panel are 0 in this data, and the rate of imprecise strategies falls to 0. However, we find that the rate of use of incomplete and inattentive strategies is actually slightly larger than in the raw dataset when we restrict to subjects who avoid imprecise strategies.

Result 6. *We find similarly strong evidence of the use of incomplete and inattentive simplification strategies among subjects who show no evidence of using imprecise strategies. This suggests cognitive or behavioral noise is likely not driving the identification and measurement of these strategies.*

Next, we have interpreted the increase in trade when subjects have to infer signal from their counterpart's choice as deriving from the use of *incomplete simplification strategies* that economize on cognitive effort by ignoring a crucial part of the reasoning required to make trading decisions. An alternative possibility is that subjects instead have a very low opinion of their counterparts' cognitive ability, and actually *believe* they trade perfectly randomly, without regard to their own signals. To examine this possibility we make use of our elicitation of subjects' beliefs about their

we are likely removing any confused subjects as well.

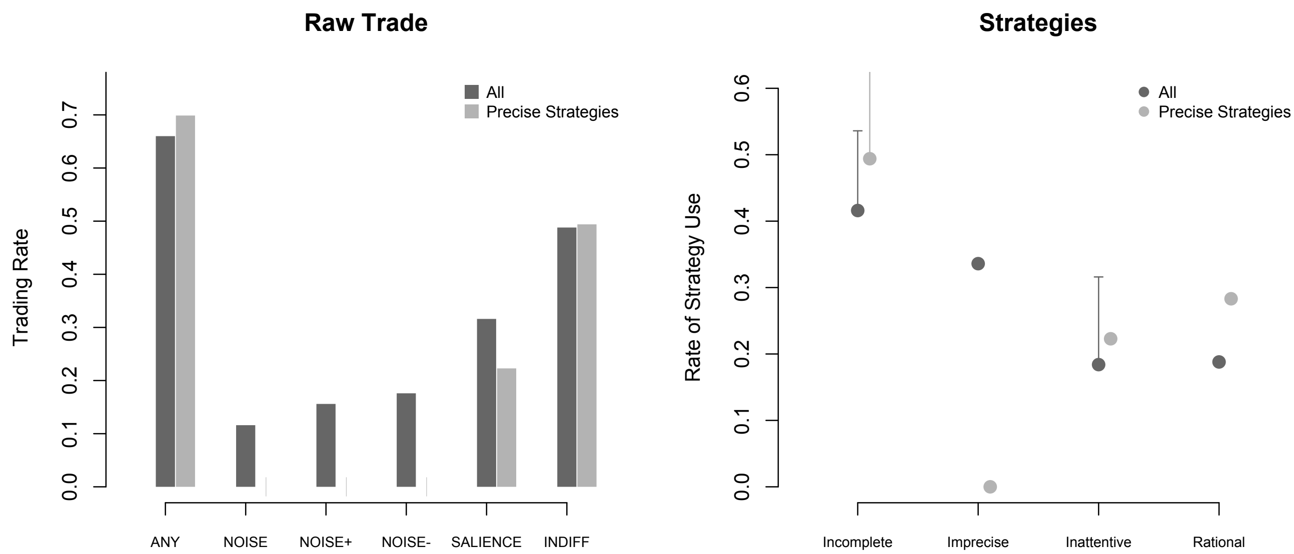

Figure 4: Rates of trade in each task (left) and estimated rates of strategy usage (right) for the full sample (dark gray) and for the subset of subjects who are not identified as using imprecise strategies (light gray). Vertical bars in the right-hand panel show the upper bound of use of the corresponding strategies.

counterpart's score in the IQ test (Raven's progressive matrix task) administered in the experiment. If failure to infer counterpart's signal is due to low opinion of counterparts' ability (rather than a simplification strategy), we would expect this behavior to disappear or shrink for subjects who believe their counterparts are of high cognitive ability.

In Figure 5, we plot data (in lighter gray) for subjects who express a belief about their counterpart's cognitive score that is in the upper quartile of such beliefs. The results show that this has no effect at all on the fraction of subjects we are able to type as "incomplete." This suggests that neglect of information in counterparts' strategies is unlikely to be built on dismissive beliefs about counterparts, and is instead driven by subjects ignoring their counterpart's signal entirely, consistent with the use of incomplete simplification strategies.

Result 7. *There is no evidence that subjects use "inattentive strategies" more when they have a lower opinion of the cognitive abilities of their counterpart. This suggests that subjects' strategies are not due to dismissive beliefs about their counterparts' sophistication, but are instead a failure to reason about the information contained in counterpart choices in the first place.*

Finally, we examine how robust the use of the simplification strategies is to subjects' own

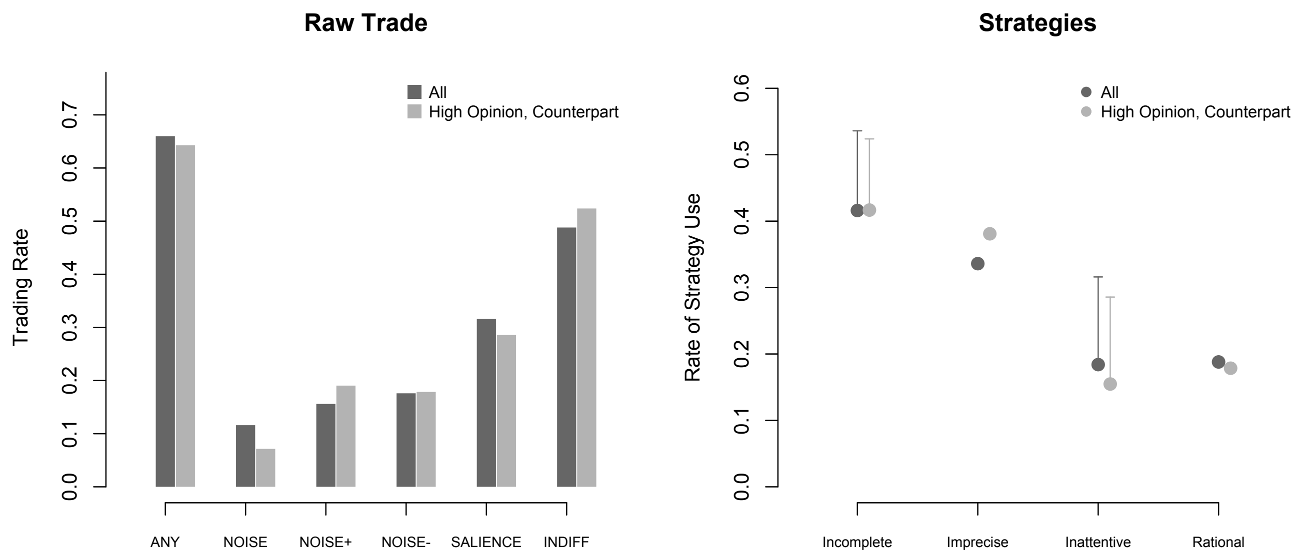

Figure 5: Rates of trade in each task (left) and estimated rates of strategy usage (right) for the full sample (dark gray) and for the subset of subjects who hold high beliefs about the cognitive ability of their counterparts (light gray). Vertical bars in the right-hand panel show the upper bound of use of the corresponding strategies.

cognitive sophistication. To do this, we examine whether evidence of use of these strategies lessens among subjects measured to have high cognitive ability in our Raven's task. In Figure 6, we plot in gray results for the subset of subjects whose Raven's score is greater or equal to the median Raven score in our sample.

The results show that subjects with high cognitive ability are slightly *more* likely than the full sample to trade in the baseline ANY task (left panel), driven by slightly higher usage of incomplete and inattentive strategies. Subjects are moderately less likely to use noisy strategies, but those subjects avoiding such strategies are measured using incomplete and inattentive strategies instead. Thus, overall, we find that these results continue to hold, broadly, even for our most cognitively capable subjects. A natural interpretation of this is that the use of the strategies we have identified may be due to deliberate, strategic efforts to limit the costs of decision-making rather than a consequence of binding constraints on cognitive resources.

Result 8. *High cognitive ability subjects are no less likely to trade than the sample as a whole, and continue to use simplification strategies at a similarly high rate.*

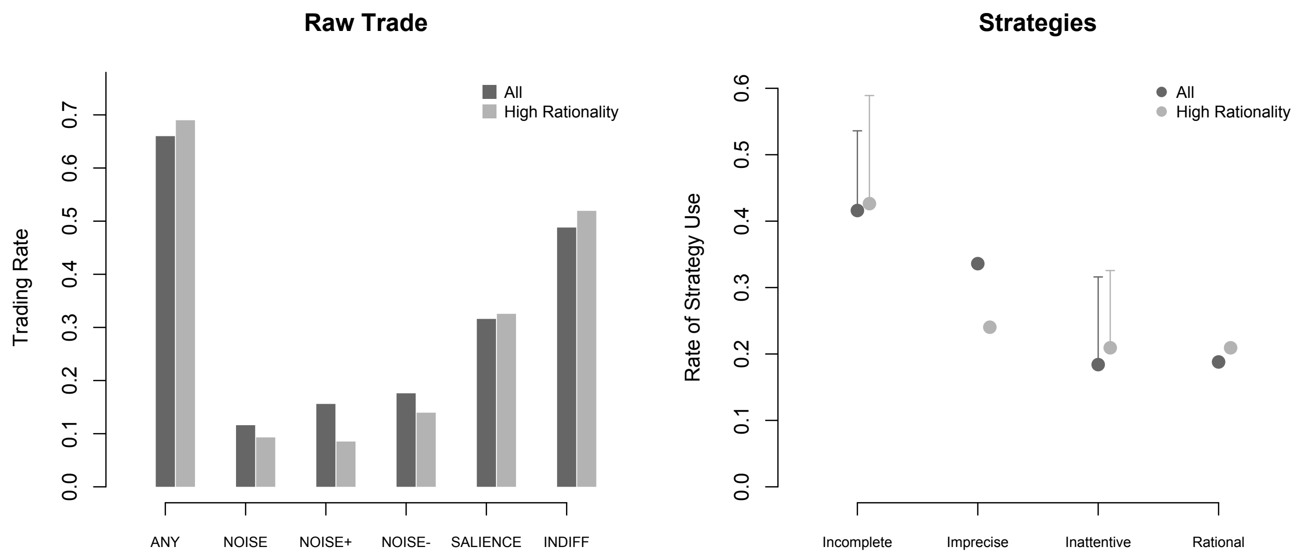

Figure 6: Rates of trade in each task (left) and estimated rates of strategy usage (right) in the full dataset.

6 The Endogenous Signals (END) Condition

In our *endogenous signals condition* (END), we repeat our entire design but change the nature of the signals, making their accuracy endogenous to subjects' own cognitive ability. In particular, subjects who score higher in the Raven's tasks administered at the end of the experiment are assigned signals with greater accuracy. The idea behind this arm of the design is to replicate a realistic feature of many real-world trading environments: traders' beliefs are often formed by their own, cognitively intensive processing of information which depends on their own cognitive skills.

We included this robustness arm of the design for two primary reasons. First, it is realistic and therefore allows us to examine the robustness of our results to an arguably more ecologically valid signal technology. Second, because signals are rooted in subjects' own cognitive ability, it potentially creates scope for trade to arise for a psychological reason that is somewhat different from the cognitive frictions that are our paper's focus. In particular, as we show below, *relative overconfidence* about one's ability relative to the ability of one's trading partner can induce that person to trade. Our aim with the design was to include an alternative psychological driver of trade (i.e., overconfident beliefs) against which to benchmark the magnitude of the cognitively-driven results from our EX condition.

6.1 Theory: Overconfidence and Trade

We formalize *relative overconfidence* by assuming that overconfident agent i believes the likelihood of his own signal being correct, $\hat{\mu}$, is higher than the likelihood of player j's signal, $\tilde{\mu}$: $\hat{\mu} > \tilde{\mu}$. As usual in the literature on overconfidence, we assume these beliefs are common knowledge. If each player i is sufficiently relatively overconfident, trading upon receiving signal $s_i = j$ becomes an equilibrium. To describe the degree of overconfidence of player i it is convenient to define the average signal-to-noise ratio of signals s_i and s_j, as perceived by the overconfident player i:

$$S\hat{N}R \equiv \frac{\hat{\mu}}{1-\hat{\mu}}$$

and

$$S\tilde{N}R \equiv \frac{\tilde{\mu}}{1-\tilde{\mu}}$$

We can then prove the following Proposition:

Proposition 5. *If* $\frac{S\hat{N}R}{S\tilde{N}R} \geq \frac{\overline{v}-\underline{v}+f}{\overline{v}-\underline{v}-f}$ *there is a Bayes-Nash equilibrium where each player* i *chooses* $a_i = T$ *if and only if* $s_i = j$.

Thus, if player i believes the signal-to-noise ratio of her information is larger than the SNR of her trading partner's information by a factor of $\frac{\overline{v}-\underline{v}+f}{\overline{v}-\underline{v}-f} > 1$, she will agree to trade. For example, if $\hat{\mu} = 0.75$, $\tilde{\mu} = 0.6$, players choose to trade for the fee level and asset payoffs considered in this experiment.

6.2 Experimental Results

The identification logic for overconfidence mirrors that of inattentive strategies (salience) in our EX condition. Because tasks c_-^+ (SALIENCE) and c_+^- (NOISE-) contains signals that are permutations of one another, overconfidence should cause an overweighting of one's own signal and therefore greater trade in the former than in the latter. Of course, we already know from the EX condition that this effect occurs already due to the use of inattentive strategies even without the influence of overconfidence. Thus the hypothesized effect of overconfidence is an *increase* in the apparent use of inattentive strategies relative to our EX condition.

Figure 7 plots results from the *END condition* in light gray; dark gray bars and points continue to represent data from our *EX condition.* Focusing first on the SALIENCE task, we find a substantial increase in the rate of trade under endogenous signals: 50% vs. 31.6%. Subtracting c_+^- from

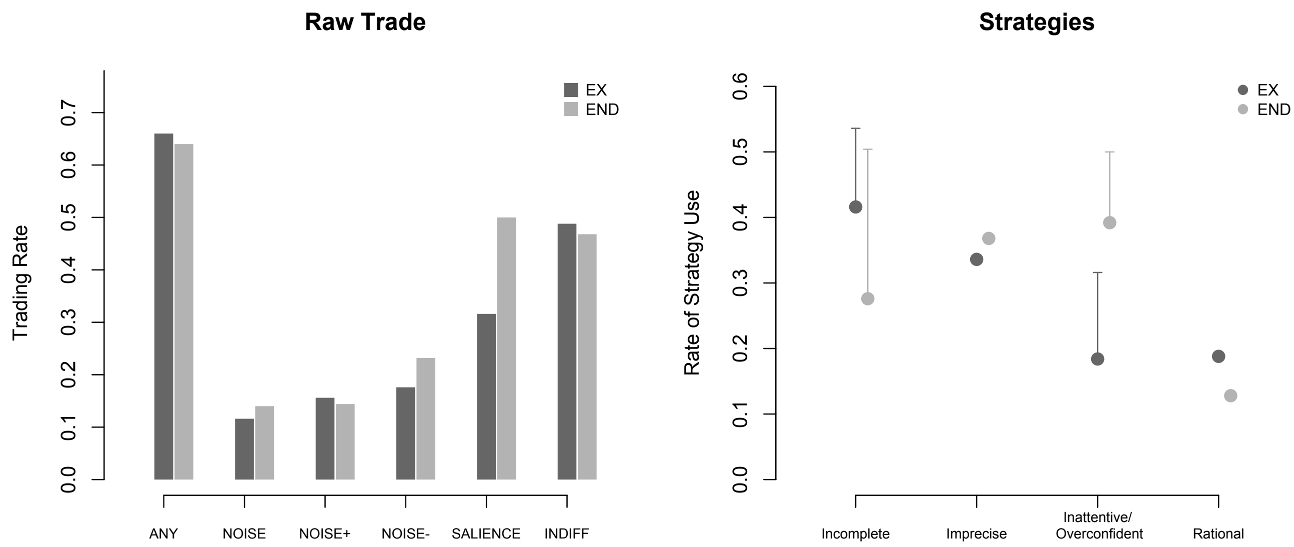

Figure 7: Rates of trade in each task (left) and estimated rates of strategy usage (right) in the full dataset, broken down by EX and END conditions. Vertical bars in the right-hand panel show the upper bound of use of the corresponding strategies.

this to get our estimate of trade attributable to overweighting of one's own signal (plotted as "inattentive" in the right panel of the Figure), we see an even larger difference: 39.2% vs. 18.4%. Thus, endogenous signals generate a *doubling* of the trade induced by overweighting own-signals, which is exactly the effect we would predict overconfidence to have.

The right hand plot also shows that our conservative measure of use of incomplete strategies falls in the endogenous signals design. This however is a mechanical effect of the increase in overweighting of own-signals (i.e. the increased use of "inattentive strategies"), since use of such strategies potentially censors measurement of incomplete strategies. As the lines extending from the "inattentive strategies" dot in that panel indicates, use of incomplete strategies may in fact be unaffected in this design.

Overall, however, as the left panel shows, this increase in trade due to overweighting of own signal has little effect on the overall rate of trade in our baseline ANY task. This is for the simple reason that overconfidence (like inattentive strategies) are substitutes for incomplete strategies: conditional on using inattentive strategies overconfidence has no further effect on trade. Thus our results for the most part are highly robust to adding this form of realism to the design.

Result 9. *Introducing endogenous, skill-based signals causes subjects to trade as though they put substantially more weight on their own signals than on their counterparts' relative to our EX condition.*

Finally, we can explicitly test whether this increase in weight on own signal in the END condition is in fact driven by relative overconfidence. Prior to making trading decisions, we asked subjects to explicitly estimate the number of Raven's questions both they and their trading partner get right, effectively allowing us to estimate $\hat{\mu}$ and $\tilde{\mu}$. Calculating the *difference*, δ between the number of Raven's questions subjects estimated they got right and the number of questions their counterpart got right, we find substantial relative overconfidence in the data. 40% of subjects believed they got more questions correct than their counterpart while only 22% believed the reverse (38% estimated them as the same).

To test whether this relative overconfidence is linked to overweighting of own signal in the data, we regress the excess trade due to weighting of own signal (SALIENCE-NOISE-) on δ, controlling for the subject's own Raven's score. We find that β, the coefficient on δ in this regression, is very small and statistically insignificant (β =0.01, p = 0.743), suggesting that there is no relationship between overconfidence and overweighting of own information in the *endogenous signals* design.

This (perhaps surprising) finding suggests that the increased overweighting we observe in END relative to EX is *not* in fact driven by overconfidence at all. Instead it seems to be an intensification of the over-weighting already observed in our *EX condition*. Or in other words, it seems that one's own endogenous signal (linked as it is to the subject's own skills) is relatively more salient than exogenously assigned signals. The consequence seems to be that this has the effect of increasing the propensity of inattentive strategies to induce trade.

Result 10. *There is no evidence that the increased effect on trade of overweighting one's own signals in the endogenous signals design is a consequence of relative overconfidence. This suggests that own-skill-based signals are relatively more salient than exogenous signals, causing them to receive more relative weight. Thus, endogenous signals induce more trade by increasing the distorting effects of inattentive strategies.*

Overconfidence is a popular explanation for trade in behavioral finance, see Daniel and Hirshleifer (2015) for an overview of this literature. Our results cast doubt on this beliefs-based explanation. Even in a simple environment like ours where the link between skill and signal accuracy is maximally transparent, and even in the presence of overconfidence over skills in the aggregate, we fail to find a statistically robust relation between overconfidence and trading at the individual level. Instead, the evidence suggests that skill-based information enhances trade by increasing the salience of private signals.[14] To conclude, this finding surprised us but turned out to reinforce the findings from our EX design. What was originally intended to measure the trade induced by a contrasting psychological force (relative overconfidence), turned out instead to identify an important (and realistic) factor that intensifies the effect of one of the simplification strategies identified in our EX data.

7 Structural Estimation

Our reduced form results suggest that incomplete simplification strategies have a particularly large effect on trade, but that under some conditions imprecise and inattentive strategies can have similarly sized effects. One of the primary costs of our reduced form approach to the data is that it ignores a potentially important effect of imprecise strategies: their capacity to induce *other subjects* to trade in equilibrium. In this section, we structurally estimate these three forces using both a model with and without equilibrium beliefs, that allows us to not only identify the relative importance of these forces in inducing trade but also the role (if any) of equilibrium forces. We use this to verify our reduced form results under the lens of our model and assess the role equilibrium forces play in driving trade in our data. Finally, we use counterfactuals from these estimates to quantitatively assess the relative importance of all of these factors in driving trade.

Model and estimation details. Our structural model builds on the formalization of the three simplification strategies introduced in Section 2.3, and extends these mechanisms to an equilibrium setting. In order to create scope for incomplete reasoning, we assume that each subject has correct beliefs about the distribution of the actions taken by her counterparts but with probability $\alpha \in [0, 1]$ fails to think about the problem completely, and thus treats her counterpart's action as *independent* of her counterpart's signal. To create scope for inattentive behavior (i.e., salience effects), each subject overweights her signal when updating posteriors as measured by the own-signal attachment parameter β. Finally, to create scope for imprecise behavior, the probability a subject proposes to trade conditional on her information depends on her beliefs about the gain from trading relative to not trading and $\gamma \geq 0$ parameterizes the sensitivity of trade to this believed gain.

[14] Previous evidence on this is limited: for example Kogan (2009) also induced skill-based signal accuracies in the lab, finding aggregate treatment effects consistent with overconfidence, but did not measure individual-level beliefs about skill.

In some of our estimates, we extend these mechanisms to an equilibrium setting by imposing the restriction that player i's beliefs about the strategy of player j are consistent with the actual play of player j. The resulting model is formally equivalent to a hybrid of the quantal response equilibrium (QRE) model (McKelvey and Palfrey, 1995; Goeree et al., 2016) and the cursed equilibrium of Eyster and Rabin (2005), with the additional feature that private signals are potentially overweighted.[15] This model features equilibrium effects: knowing that one's counterpart will agree to trade with some probability regardless of her signal increases the conditional probability that a trade is profitable given that one's counterpart agrees to trade, $\pi_{i,k}$. This, in turn, encourages more trade from i, generating an amplified equilibrium effect.

Using the standard logit specification to model the generation of noise, the probability player i chooses to trade conditional on signal (or signal profile) k, $\sigma_{i,k}$, is given by

$$\sigma_{i,k} \equiv \Pr(a_i = T \mid s_i = k) = \frac{e^{\Delta_{i,k}/\gamma}}{1 + e^{\Delta_{i,k}/\gamma}}. \tag{3}$$

where, as in Section 2.3, $\Delta_{i,k}$ is the perceived gain from trading, given by:

$$\Delta_{i,k} = \{\tau_{i,k}\left[2(\overline{v} - \underline{v})(\pi_{i,k} - 0.5) - f\right]\} \tag{4}$$

As before, $\tau_{i,k}$ is player i's conditional belief that her counterpart will trade and $\pi_{i,k}$ is the conditional probability she attaches to the state being j (i.e. that it is advantageous for her to trade).[16]

The specification of beliefs $\tau_{i,k}$ and $\pi_{i,k}$ depends on whether signals are private or public. The model also allows for subjective beliefs about signal accuracies, but we focus on our main exogenous information treatments for now.

In private information treatments, signals are private so that beliefs $\tau_{i,k}$ and $\pi_{i,k}$ can be described by the following expressions:

$$\tau_{i,k} = \begin{cases} \sigma_{j,i}\left[\frac{\mu^\beta}{\mu^\beta+(1-\mu)^\beta}(1-\mu) + \frac{(1-\mu)^\beta}{\mu^\beta+(1-\mu)^\beta}\mu\right] + \sigma_{j,j}\left[\frac{\mu^\beta}{\mu^\beta+(1-\mu)^\beta}\mu + \frac{(1-\mu)^\beta}{\mu^\beta+(1-\mu)^\beta}(1-\mu)\right] & \text{if } k = j \\ \sigma_{j,i}\left[\frac{\mu^\beta}{\mu^\beta+(1-\mu)^\beta}\mu + \frac{(1-\mu)^\beta}{\mu^\beta+(1-\mu)^\beta}(1-\mu)\right] + \sigma_{j,j}\left[\frac{\mu^\beta}{\mu^\beta+(1-\mu)^\beta}(1-\mu) + \frac{(1-\mu)^\beta}{\mu^\beta+(1-\mu)^\beta}\mu\right] & \text{if } k = i \end{cases}$$

$$\pi_{i,k} = \begin{cases} \alpha\frac{\mu^\beta}{\mu^\beta+(1-\mu)^\beta} + (1-\alpha)\frac{\mu^\beta[\sigma_{j,i}(1-\mu)+\sigma_{j,j}\mu]}{\mu^\beta[\sigma_{j,i}(1-\mu)+\sigma_{j,j}\mu]+(1-\mu)^\beta[\sigma_{j,i}\mu+\sigma_{j,j}(1-\mu)]} & \text{if } k = j \\ \alpha\frac{(1-\mu)^\beta}{\mu^\beta+(1-\mu)^\beta} + (1-\alpha)\frac{(1-\mu)^\beta[\sigma_{j,i}(1-\mu)+\sigma_{j,j}\mu]}{(1-\mu)^\beta[\sigma_{j,i}(1-\mu)+\sigma_{j,j}\mu]+\mu^\beta[\sigma_{j,i}\mu+\sigma_{j,j}(1-\mu)]} & \text{if } k = i \end{cases}$$

[15] Carrillo and Palfrey (2009) and Camerer et al. (2016) estimate structural models that combine QRE and cursed beliefs on other games. The distinctive feature of our approach is that we use this structural model to evaluate the magnitudes of the underlying cognitive frictions and the way they are amplified in equilibrium.

[16] The model is symmetric and therefore the expressions for trade probabilities, payoffs and beliefs are the same for the two players (exchanging the signals as necessary). We keep the player's subscript in the expressions to make them easier to interpret.

Beliefs are shaped by behavioral parameters α, β and, indirectly, γ; in our main exogenous information treatments, $\mu = 0.7$ by design.[17]

In tasks in which signals are public, player i's beliefs and actions are conditioned on a public signal profile. For instance, $\sigma_{i,(k,l)}$ now represents the probability of action $a_i = T$ conditional on the public signal profile $\langle s_i = k, s_j = l \rangle$. Player i's beliefs are given by the following expressions:

$$\tau_{i,(k,l)} = \sigma_{j,(l,k)}$$

$$\pi_{i,(k,l)} = \begin{cases} \frac{(1-\mu)^\beta(1-\mu)}{(1-\mu)^\beta(1-\mu)+\mu^\beta\mu} & \text{if } k=i, l=i \\ \frac{\mu^\beta(1-\mu)}{\mu^\beta(1-\mu)+(1-\mu)^\beta\mu} & \text{if } k=j, l=i \\ \frac{(1-\mu)^\beta\mu}{\mu^\beta(1-\mu)+(1-\mu)^\beta\mu} & \text{if } k=i, l=j \\ \frac{\mu^\beta\mu}{(1-\mu)^\beta(1-\mu)+\mu^\beta\mu} & \text{if } k=j, l=j \end{cases}$$

In these tasks, posterior beliefs are unaffected by limited contingent thinking by design: the parameter α does not enter the expressions in this case. Again, in our main exogenous information treatments, $\mu = 0.7$ by design.

A similar model can be estimated using data from the endogenous information treatments. To do this, we modify the expressions above by replacing the signal accuracy parameter μ with subject i's beliefs about accuracies. In turn, to calibrate a subject's perceived accuracy we use our data on Raven's test forecasts. However, we also allow for some disconnect between perceived signal accuracies and Raven's test forecasts, since the reduced form evidence presented in subsection 6 shows that belief-based measures of overconfidence do not explain trading decisions. We assume that subject i perceives the accuracy of her own signal to be given by: $\delta\mu + (1-\delta)\hat{\mu}_i$, where $\mu = 0.7$ is the exogenous signal accuracy in our main treatments and $\hat{\mu}_i$ is the signal accuracy computed using subject i's forecast of her own Raven's test score. Similarly, we assume subject i perceives the accuracy of her counterpart's signal to equal: $\delta\mu + (1-\delta)\tilde{\mu}_i$. The parameter $\delta \in [0, 1]$ measures the disconnect between beliefs about skills and perceived signal accuracies. When $\delta = 0$, the subjects correctly perceive signal accuracies in the endogenous information treatment to be determined by skill. When $\delta = 1$, subjects ignore this link and estimate signal accuracies to be exogenous (and equal to 0.7 as in the exogenous information treatment).

While so far we have assumed that other parameters are the same in the endogenous information treatment as in the exogenous information treatment, in a version of our model we allow the

[17] The limited contingent thinking parameter α does not enter the expression for $\tau_{i,k}$ because, as in a cursed equilibrium, players have correct beliefs about the distribution of other players' actions (while neglecting the relation between actions and private information).

attention parameter β to depend on whether signals are endogenous or exogenous. This possibility is inspired by the fact that our reduced-form measure of the effect of inattentive strategies is larger when information is endogenous (see subsection 6). To measure this mechanism in our structural model, we estimate a distinct attention parameter for each treatment, labelled β_{END} and β_{EX} respectively.

Finally, and importantly, we also estimate a *non-equilibrium model* using data from the EX condition. The previous models assume each subject believes that her counterpart's trading decision is driven by the same incentives and cognitive frictions as her own. Formally, player i correctly anticipates the conditional probability of player j choosing to trade, which is in turn determined by the model. We are interested in understanding whether this assumption actually improves the explanatory power of our behavioral theory. In order to answer this question, we estimate a non-equilibrium model where each subject's beliefs about her counterpart's strategies are constant and independent of the other model parameters. We assume that each player i believes player $j \neq i$ will trade whenever player j's signal is favourable, that is: $\sigma_{j,i} = 1, \sigma_{j,j} = 0, \sigma_{j,(i,k)} = 1 \forall k, \sigma_{j,(j,k)} = 0 \forall k$. We view this as a plausible alternative assumption about strategically naive beliefs.

Estimation results. Structural estimation results are reported in Table 1. In column (1), we show the results for the model estimated on data from treatments with exogenous information only. We estimate $\alpha = 1$ suggesting that subjects use incomplete strategies and on average completely ignore information in others' choices. Importantly, our estimate of incomplete strategies is not confounded by the equilibrium effect of noisy decisions, as our structural model controls for it. In fact, incomplete strategies cause a significant loss in information since in actuality trading partners are four times more likely to trade on their signal than against it in the private information tasks. We find strong evidence that subjects overweight their own salient signal: the estimate of β is significantly larger than unity (p-value< 0.001). Finally, our estimate of γ show that there is considerable noise in decisions, since the parameter is significantly greater than zero (p-value< 0.001), indicating the use of imprecise strategies. Thus, our estimates qualitatively match and give us a sharper quantitative sense of our reduced form results from the EX treatment.

When we re-estimate the model using both data from exogenous and endogenous information treatments (EX and END), we find very similar results. The estimates reported in column (2) show subjects use incomplete strategies ($\alpha = 1$), inattentive strategies ($\beta > 1$) and imprecise strategies ($\gamma > 0$). Moreover, our estimate of the δ parameter suggests that subjects' perceived signal accuracies in the END treatments are mostly anchored to the signal accuracy of the EX treatments

		EX and END pooled		
	EX data only	Constant attention	ENDO-specific attention	EX non-equilibrium
	(1)	(2)	(3)	(4)
α	1.00*	1.00**	1.00**	0.80**
	(0.40)	(0.32)	(0.32)	(0.25)
β	1.36***	1.53***		1.17***
	(0.11)	(0.80)		(0.10)
γ	1.69***	1.93***	1.91***	3.15***
	(0.14)	(0.12)	(0.12)	(0.25)
δ		0.75***	0.88***	
		(0.15)	(0.16)	
β_{EX}			1.44***	
			(0.10)	
β_{END}			1.67***	
			(0.14)	

Note: Estimated parameters: MLE standard errors in parentheses. *p<0.05; **p<0.01; ***p<0.001.

Table 1: Structural model parameters.

and only weakly related to their beliefs about skills (the weight on the former is estimated at 75%, while the weight on the latter is 25%). In column (3), we estimate a version of the model that allows the attention parameter β to depend on whether the signal accuracy is exogenous or endogenous. We find that subjects behave as if they overweight their own signal to a higher degree when signals are endogenous than exogenous ($\beta_{END} > \beta_{EX}$), although this difference is not statistically significant (using a likelihood ratio test, p-value= 0.16). Thus, our structural estimates seem broadly consistent with our reduced-form evidence, showing that trading decisions in the endogenous information treatments are not mainly driven by overconfidence, but possibly by an increase in the salience of a player's own signal.

The results from the non-equilibrium model are presented in column (4). This model also shows a significant effect of incomplete strategies, inattentive strategies and imprecise strategies. Relative to the equilibrium model, estimates from the non-equilibrium model suggest slightly lower degrees of incompleteness and inattentiveness, but a higher degree of noise in the strategies of subjects. We evaluate the relative performance of the equilibrium vs. non-equilibrium model next.

The role of equilibrium beliefs. How important are strategic (equilibrium) forces in driving trade? To find out we ask whether assuming equilibrium beliefs improve the explanatory power of the model relative to a simpler non-strategic accounting of cognitive frictions? We measure the explanatory power of the model using the log-likelihood at the estimated value of parameters and we report it for both the equilibrium and non-equilibrium models in Table 2. The log-likelihood of the equilibrium model is -847, while the log-likelihood of the non-equilibrium model is -893. Because the log-likelihood of the equilibrium model is higher, this model performs better. To evaluate whether such difference in performance is statistically significant we use a model selection test for strictly non-nested models (Vuong, 1989). The test rejects the null hypotheses that the two models are equally good with a p-value below 0.001. Thus, the equilibrium model performs significantly better than the non-equilibrium model.

	Equilibrium model	Non-equilibrium model
Actual data	-847	-893
Synthetic data average	-998	-1009

Table 2: Log-likelihood of equilibrium vs. non-equilibrium models

One natural concern with this conclusion is that the equilibrium model may perform better simply because it is more flexible. Although the two models have the same number of parameters, beliefs are fixed in the non-equilibrium model while they may adjust to increase model fit in the equilibrium model (although this adjustment is disciplined by the equilibrium requirement). To evaluate this concern, we use an approach similar to Fudenberg et al. (2023). We simulate 100 synthetic datasets, each with the same structure of the data from our exogenous information treatment. In each dataset, we draw a random distribution of trading decisions: for each possible signal (or signal profile), each subject trades with probability p. In turn p is drawn from a uniform distribution between 0 and 1 across datasets. We then estimate the equilibrium and non-equilibrium models on each dataset and compute the log-likelihood scores. Finally, we take the average scores across the 100 datasets.

The results are shown in Table 2. We find that the average log-likelihood scores across synthetic datasets are -998 for the equilibrium model and -1009 for the non-equilibrium model. This shows that the equilibrium model results in better fit in general, because of its flexibility. However the average gain in fitness (log-likelihood score) is only 11 across the synthetic datasets, while it is 46 in our own data. This suggests that the higher fitness of the equilibrium model in our data goes above and beyond the advantage derived from additional flexibility. We interpret this as evidence that equilibrium forces are likely an important factor in explaining trading decisions, alongside the cognitive frictions themselves.

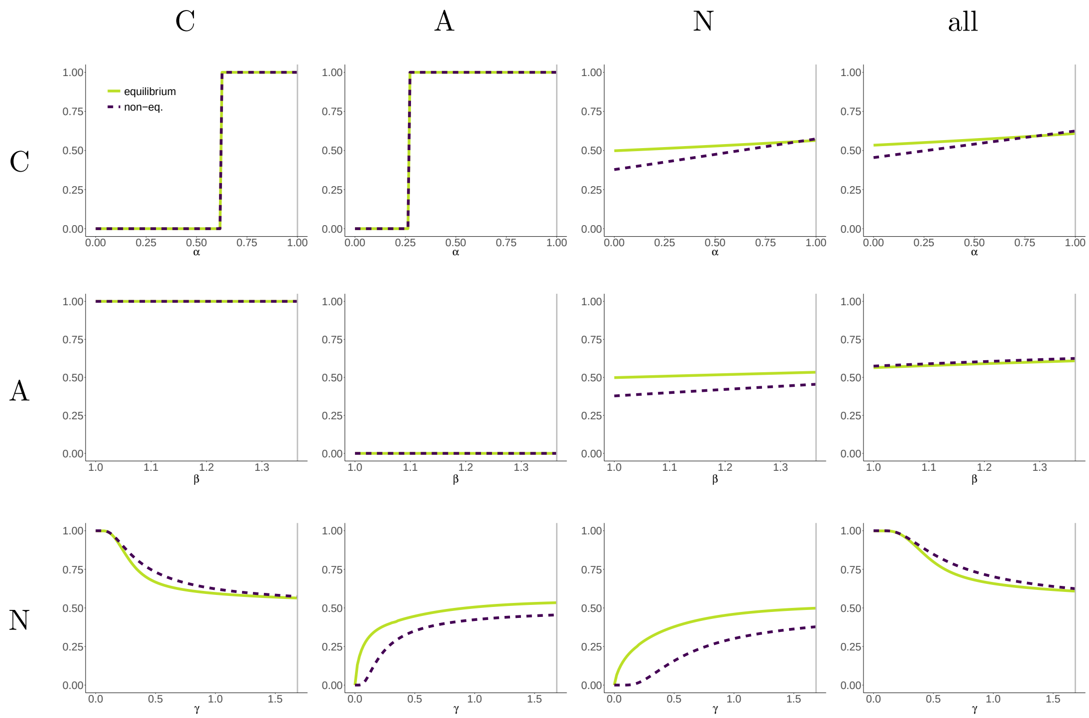

Figure 8: Trade Probability Predicted by the Estimated Model

Relative magnitudes. Finally, we use our estimated model to evaluate the magnitudes of the cognitive frictions we've examined, including their interactions and equilibrium effects. To do this, we compute the effect each simplification strategy is predicted by the estimated model to have (as a function of our structural measure of the intensity of the relevant parameter) on the probability that player i chooses to trade on her signal in the private information game. For computing these predictions, we use our estimates from data on the exogenous information (EX) condition.

We plot the results in Figure 8. We vary one cognitive friction at a time in each row of the figure: incomplete contingent reasoning (C), inattention (A) and imprecision (N). The remaining friction parameters are either set to zero or to their estimated value. Each of the four columns of Figure 8 shows which *other frictions* are operative. For instance, the top-left panel (C+C) plots the effect of the incomplete strategy parameter (α) in the counterfactual where the *only* friction is incomplete strategies. On the other hand, the second panel from the left in the top row (C+A) plots the effect of the incomplete strategy parameter (α) in the counterfactual where there is *also* scope for inattentive strategies. In each panel, the grey vertical line on the right side of each panel shows the estimated value of the parameter (we are interested in gauging the effect of reducing and potentially eliminating a given cognitive friction, so we only plot points to the left of the grey line).

As discussed above, cognitive imprecision can have important equilibrium effects, since when a player anticipates others will sometimes trade against their information, this affects her belief that a trade is profitable and thus the rate of trade. To isolate the equilibrium effect, in each case we also analyze a counterfactual (plotted with a dashed line) for the non-equilibrium model (e.g., where a player's counterpart is willing to trade if and only if her signal suggests it is worth doing so ($\sigma_{ji} = 1, \sigma_{jj} = 0$), and therefore is unaffected by noise). We keep all other parameters constant to isolate the effect of equilibrium beliefs alone. The difference in player i's rate of trade between the equilibrium model (solid line) and the non-equilibrium (dashed line) counterfactual captures the effect of player i anticipating that player j's choices are as noisy as his own.

The results show that, given the intensities we estimate, some of the cognitive frictions are powerful enough to generate high trade rates even in isolation from other frictions. Given our estimates, incomplete strategies alone can sustain 100% trade rates (panel C+C). Imprecise strategies alone can generate a trade rate around 50% (panel N+N). Inattentive strategies cannot however generate positive trade in isolation (panel A+A). Instead, inattentive strategies complement the other channels: adding scope for inattentive strategies increases trade rates in the presence of imprecise strategies (panel N+A) and it lowers the minimum degree of strategy incompleteness (alpha) required to sustain trade in equilibrium (panel C+A). When all frictions are present, removing one at a time is not effective at moderating trade rates. For example, removing scope for incomplete strategies reduces trade rates by around 6 percentage points (panel C+all). Removing noise would actually increase trade rates (panel N+all), since choices would become more responsive to the misperceived positive gains from trading. This suggests that our three channels act as substitutes, diminishing the marginal effect of one another on trade. Because of this, in order to significantly reduce trade, all three of these channels must be removed simultaneously.

Comparing dashed and solid lines in these panels, equilibrium effects influence trade, but only in some of the counterfactuals. For instance, equilibrium effects have little impact on trade when we shut down imprecise strategies (panels C+C,C+A, A+C, C+C and other panels for low values of gamma). This highlights the key role of noise in generating equilibrium effects. Equilibrium effects have the strongest positive effect on trade rates when strategies are imprecise but complete (panels N+A, N+N, A+N). Under incomplete strategies, equilibrium effects tend to be smaller (e.g. panel N+C or for high values of alpha in panels C+N and C+all). The reason for this is that incomplete strategies ignore the counterpart's strategies by construction and so make behavior less responsive to strategic beliefs already. Interestingly, however, incomplete strategies have a larger impact on trade rates when beliefs are not in equilibrium (panels C+N and C+all). For instance, removing scope for incomplete strategies reduces trade rates by around 20 percentage points in the non-equilibrium counterfactual, as opposed to the mere 6 percentage points effect under equilibrium beliefs (panel C+all). The reason for this is that incomplete contingent thinking and noise act as substitutes. When trading partners adopt very imprecise strategies, and thus their trading decisions are only weakly based on their private information, failing to think about the information content of a counterpart's willingness to trade is a small mistake and does not generate additional willingness to trade.

Summarizing the results of our structural analysis:

Result 11. *Structural estimates show:*

(i) incomplete, imprecise and inattentive strategies all have a significant effect on trade,

(ii) when combined, these strategies generate a trade rate above 60%,

(iii) incomplete and imprecise strategies have large effects on the rate of trade in isolation,

(iv) the simplification strategies are substitutes, diminishing the effect of one another on trade,

(v) equilibrium beliefs are important in quantitatively explaining behavior and

(vi) in equilibrium, the effect of imprecise strategies is magnified, while the effect of incomplete strategies and inattentive strategies is attenuated.

8 Conclusion

A recent literature in "cognitive economics" has identified a number of domain-general cognitive frictions that lead to systematic departures from standard economic predictions. This literature has shown that a relatively small set of cognitive frictions can help to explain a large number of famous anomalies documented in behavioral economics. However this literature has largely been focused on understanding individual choice. We take some early steps to understand how these same frictions influence the behavior of markets. To do this, we examine the role several of the most important frictions identified by this literature play in one of the key mysteries in market microstructure: the tendency of people to engage in speculative trade despite the severe adverse selection problems described by "no-trade" theorems.

Using a novel experimental design, we find very high rates of trade – in violation of no-trade theorems – even in (arguably) the simplest possible markets. We show that this trade is especially driven by the use of "incomplete" cognitive strategies that simplify the exchange problem by ignoring the information implicit in counterparts' willingness to trade. However we also find that trade is driven by the use of "imprecise" strategies that noisily implement decisions and "inattentive" strategies that imperfectly weight sources of information. By contrast, we find little evidence that relatively overconfident beliefs – a popular alternative psychological driver of trade – generates trade in our setting. Finally we show that some of these frictions are meaningfully amplified by equilibrium effects on beliefs, driving some of the trade we observe. We believe our finding that a small collection of cognitive frictions fundamentally shape trade in even our simple markets suggest that they may have an important role in helping us to understand behavior in a far broader range of market settings.

References

Angrisani, Marco, Antonio Guarino, Steffen Huck, and Nathan C Larson, "No-trade in the laboratory," *The BE Journal of Theoretical Economics*, 2008, *11* (1).

Augenblick, Ned, Eben Lazarus, and Michael Thaler, "Overinference from weak signals and underinference from strong signals," *The Quarterly Journal of Economics*, 2025, *140* (1), 335–401.

Azrieli, Yaron, Christopher P Chambers, and Paul J Healy, "Incentives in experiments: A theoretical analysis," *Journal of Political Economy*, 2018, *126* (4), 1472–1503.

Ba, Cuimin, J Aislinn Bohren, and Alex Imas, "Over-and Underreaction to Information: Belief Updating with Cognitive Constraints," Working Paper, Department of Economics, University of Pennsylvania 2025.

Barber, Brad M and Terrance Odean, "All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors," *The Review of Financial Studies*, 2008, *21* (2), 785–818.

Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer, "How people use statistics," *Review of Economic Studies*, 2026, *93* (1), 250–285.

__ , **John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer**, "Memory and probability," *The Quarterly Journal of Economics*, 2023, *138* (1), 265–311.

__ , **Nicola Gennaioli, and Andrei Shleifer**, "Salience in experimental tests of the endowment effect," *American Economic Review*, 2012, *102* (3), 47–52.

__ , __ , **and** __ , "Salience theory of choice under risk," *The Quarterly journal of economics*, 2012, *127* (3), 1243–1285.

__ , __ , **and** __ , "Salience and consumer choice," *Journal of Political Economy*, 2013, *121* (5), 803–843.

__ , __ , **and** __ , "Salience," *Annual Review of Economics*, 2022, *14* (1), 521–544.

Brocas, Isabelle, Juan D Carrillo, Stephanie W Wang, and Colin F Camerer, "Imperfect choice or imperfect attention? Understanding strategic thinking in private information games," *The Review of Economic Studies*, 2014, p. rdu001.

Burks, Stephen V, Jeffrey P Carpenter, Lorenz Goette, and Aldo Rustichini, "Overconfidence and social signalling," *Review of Economic Studies*, 2013, *80* (3), 949–983.

Camerer, Colin, Salvatore Nunnari, and Thomas R. Palfrey, "Quantal response and nonequilibrium beliefs explain overbidding in maximum-value auctions," *Games and Economic Behavior*, 2016, *98*, 243 – 263.

Carrillo, Juan D and Thomas R Palfrey, "The compromise game: two-sided adverse selection in the laboratory," *American Economic Journal: Microeconomics*, 2009, *1* (1), 151–181.

__ **and** __ , "No trade," *Games and Economic Behavior*, 2011, *71* (1), 66–87.

Charness, Gary and Dan Levin, "The origin of the winner's curse: a laboratory study," *American Economic Journal: Microeconomics*, 2009, *1* (1), 207–236.

Chen, Daniel L, Martin Schonger, and Chris Wickens, "oTree: An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Conlon, John J, Malavika Mani, Gautam Rao, Matthew W Ridley, and Frank Schilbach, "Not learning from others," Working Paper, National Bureau of Economic Research 2022.

Daniel, Kent and David Hirshleifer, "Overconfident investors, predictable returns, and excessive trading," *The Journal of Economic Perspectives*, 2015, *29* (4), 61–87.

Daniel, Kent D, David Hirshleifer, and Avanidhar Subrahmanyam, "Overconfidence, arbitrage, and equilibrium asset pricing," *The Journal of Finance*, 2001, *56* (3), 921–965.

Danz, David, Lise Vesterlund, and Alistair J Wilson, "Belief elicitation and behavioral incentive compatibility," *American Economic Review*, 2022, *112* (9), 2851–2883.

Dow, James and Gary Gorton, "Noise trading, delegated portfolio management, and economic welfare," *Journal of Political Economy*, 1997, *105* (5), 1024–1050.

Ebert, Jane EJ and Drazen Prelec, "The fragility of time: Time-insensitivity and valuation of the near and far future," *Management science*, 2007, *53* (9), 1423–1438.

Enke, Benjamin, "The Cognitive Turn in Behavioral Economics," 2026. Working paper.

_ and Thomas Graeber, "Cognitive Uncertainty," *The Quarterly Journal of Economics*, 2023.

_ , _ , and Ryan Oprea, "Complexity and time," *Journal of the European Economic Association*, 2025, *23* (5), 1838–1867.

_ , _ , _ , and Jeffrey Yang, "Behavioral attenuation," Working Paper, National Bureau of Economic Research 2024.

Esponda, Ignacio and Emanuel Vespa, "Hypothetical thinking and information extraction in the laboratory," *American Economic Journal: Microeconomics*, 2014, *6* (4), 180–202.

_ and _ , "Contingent thinking and the sure-thing principle: Revisiting classic anomalies in the laboratory," *Review of Economic Studies*, 2023, p. rdad102.

Eyster, Erik and Matthew Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

_ , _ , and Dimitri Vayanos, "Financial markets where traders neglect the informational content of prices," *The Journal of Finance*, 2019, *74* (1), 371–399.

Frederick, Shane and Daniel Kahneman, "Representativeness revisited: Attribute substitution in intuitive judgment," in Thomas Gilovich, Dale Griffin, and Daniel Kahneman, eds., *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge University Press, 2002, pp. 49–81.

French, Kenneth R, "Presidential address: The cost of active investing," *The Journal of Finance*, 2008, *63* (4), 1537–1573.

Frydman, Cary and Lawrence J Jin, "Efficient coding and risky choice," *The Quarterly Journal of Economics*, 2022, *137* (1), 161–213.

_ and _ , "On the source and instability of probability weighting," Working Paper, National Bureau of Economic Research 2023.

Fudenberg, Drew, Wayne Gao, and Annie Liang, "How flexible is that functional form? Quantifying the restrictiveness of theories," *Review of Economics and Statistics*, 2023, pp. 1–50.

Gabaix, Xavier, "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

Gennaioli, Nicola and Andrei Shleifer, "What comes to mind," *The Quarterly journal of economics*, 2010, *125* (4), 1399–1433.

Glosten, Lawrence R and Paul R Milgrom, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, 1985, *14* (1), 71–100.

Goeree, Jacob K, Charles A Holt, and Thomas R Palfrey, *Quantal Response Equilibrium: A Stochastic Theory of Games*, Princeton University Press, 2016.

_ , Thomas R Palfrey, Brian W Rogers, and Richard D McKelvey, "Self-correcting information cascades," *The Review of Economic Studies*, 2007, *74* (3), 733–762.

Hartzmark, Samuel M, "The worst, the best, ignoring all the rest: The rank effect and trading behavior," *The Review of Financial Studies*, 2015, *28* (4), 1024–1059.

Ivanov, Asen, Dan Levin, and Muriel Niederle, "Can relaxation of beliefs rationalize the winner's curse?: an experimental study," *Econometrica*, 2010, *78* (4), 1435–1452.

Khaw, Mel Win, Ziang Li, and Michael Woodford, "Cognitive imprecision and small-stakes risk aversion," *The Review of Economic Studies*, 2021, *88* (4), 1979–2013.

Kogan, Shimon, "Distinguishing the effect of overconfidence from rational best-response on information aggregation," *The Review of Financial Studies*, 2009, *22* (5), 1889–1914.

Kyle, Albert S, "Continuous auctions and insider trading," *Econometrica: Journal of the Econometric Society*, 1985, pp. 1315–1335.

Loewenstein, George and Zachary Wojtowicz, "The Economics of Attention," *Journal of Economic Literature*, September 2025, *63* (3), 1038–1089.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa, "Failures in contingent reasoning: The role of uncertainty," *American Economic Review*, 2019, *109* (10), 3437–3474.

Matějka, Filip and Alisdair McKay, "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 2015, *105* (1), 272–298.

McKelvey, Richard D and Thomas R Palfrey, "Quantal Response Equilibria for Normal Form Games," *"Games and Economic Behavior "*, 1995, *10*, 6–38.

Milgrom, Paul and Nancy Stokey, "Information, trade and common knowledge," *Journal of Economic Theory*, 1982, *26* (1), 17–27.

Nagel, Lea, Muriel Niederle, and Emanuel Vespa, "Decomposing the winner's curse," Working Paper 2025.

Ngangoué, M. Kathleen and Georg Weizsäcker, "Learning from Unrealized versus Realized Prices," *American Economic Journal: Microeconomics*, May 2021, *13* (2), 174–201.

Niederle, Muriel and Emanuel Vespa, "Cognitive limitations: Failures of contingent thinking," *Annual Review of Economics*, 2023, *15* (1), 307–328.

Nöth, Markus and Martin Weber, "Information aggregation with random ordering: Cascades and overconfidence," *The Economic Journal*, 2003, *113* (484), 166–189.

Odean, Terrance, "Volume, volatility, price, and profit when all traders are above average," *The Journal of Finance*, 1998, *53* (6), 1887–1934.

Oprea, Ryan, "Decisions under risk are decisions under complexity," *American Economic Review*, 2024, *114* (12), 3789–3811.

_ **and Sevgi Yuksel**, "Social exchange of motivated beliefs," *Journal of the European Economic Association*, 2022, *20* (2), 667–699.

Plott, Charles R and Shyam Sunder, "Rational expectations and the aggregation of diverse information in laboratory security markets," *Econometrica: Journal of the Econometric Society*, 1988, pp. 1085–1118.

Polanía, Rafael, Michael Woodford, and Christian C Ruff, "Efficient coding of subjective value," *Nature neuroscience*, 2019, *22* (1), 134–142.

Ross, Stephen A, "Discussion: Intertemporal asset pricing," in Sudipto Bhattacharya and George M Constantinides, eds., *Theory of valuation: Frontiers of modern financial theory*, Vol. 1, Rowman & Littlefield, 1989, pp. 85–96.

Scheinkman, Jose A and Wei Xiong, "Overconfidence and speculative bubbles," *Journal of political Economy*, 2003, *111* (6), 1183–1220.

Shleifer, Andrei and Lawrence H Summers, "The noise trader approach to finance," *Journal of Economic Perspectives*, 1990, *4* (2), 19–33.

Sterling, Peter and Simon Laughlin, *Principles of neural design*, MIT press, 2015.

Tirole, Jean, "On the possibility of speculation under rational expectations," *Econometrica: Journal of the Econometric Society*, 1982, pp. 1163–1181.

Vuong, Quang H, "Likelihood ratio tests for model selection and non-nested hypotheses," *Econometrica: journal of the Econometric Society*, 1989, pp. 307–333.

Weizsäcker, Georg, "Do we follow others when we should? A simple test of rational expectations," *American Economic Review*, 2010, *100* (5), 2340–2360.

Woodford, Michael, "Stochastic choice: An optimizing neuroeconomic model," *American Economic Review*, 2014, *104* (5), 495–500.

_ , "Modeling imprecision in perception, valuation, and choice," *Annual Review of Economics*, 2020, *12* (1), 579–601.

Appendix

A Proofs

A.1 Proof of Proposition 1

Proof. By contradiction, suppose there is a Bayes-Nash equilibrium in which both players choose T for some signal. Specifically, suppose player i chooses $a_i = T$ if and only if $s_i \in S_i$, where $S_i \subseteq \{1,2\}$ is non-empty, and player $j \neq i$ chooses $a_j = T$ if and only if $s_j \in S_j$, where $S_j \subseteq \{1,2\}$ is non-empty.

In equilibrium, for every $k \in S_i$, player i must weakly prefer choosing T to choosing N, so $\Delta_{i,k} \geq 0$. Hence

$$\tau_{i,k}\left\{[\pi_{i,k}u_i(\overline{v}-f)+(1-\pi_{i,k})u_i(\underline{v}-f)]-[\pi_{i,k}u_i(\underline{v})+(1-\pi_{i,k})u_i(\overline{v})]\right\} \geq 0.$$

Since $k \in S_i$ and S_j is non-empty, we must have $\tau_{i,k} = \Pr(a_j = T \mid s_i = k) > 0$. Also, because $f > 0$ and u_i is strictly increasing,

$$u_i(\overline{v}) > u_i(\overline{v}-f) \qquad \text{and} \qquad u_i(\underline{v}) > u_i(\underline{v}-f).$$

Therefore, if the above inequality holds, it must also be that

$$[\pi_{i,k}u_i(\overline{v})+(1-\pi_{i,k})u_i(\underline{v})]-[\pi_{i,k}u_i(\underline{v})+(1-\pi_{i,k})u_i(\overline{v})] > 0.$$

Rearranging,

$$2\big[u_i(\overline{v})-u_i(\underline{v})\big](\pi_{i,k}-1/2) > 0.$$

Since $u_i(\overline{v}) > u_i(\underline{v})$, it follows that

$$\pi_{i,k} > 1/2.$$

Using the definition of $\pi_{i,k}$ and the fact that $a_j = T$ iff $s_j \in S_j$,

$$\Pr(\omega = j \mid s_i = k,\ s_j \in S_j) > 1/2 \qquad \text{for every } k \in S_i.$$

Now condition on the event $s_i \in S_i$. Since

$$\Pr(\omega = j \mid s_i \in S_i,\ s_j \in S_j)$$

is a weighted average of

$$\Pr(\omega = j \mid s_i = k,\ s_j \in S_j), \qquad k \in S_i,$$

and each of these terms is strictly greater than 1/2, we obtain

$$\Pr(\omega = j \mid s_i \in S_i,\ s_j \in S_j) > 1/2.$$

By the same argument, for player j we obtain

$$\Pr(\omega = i \mid s_i \in S_i,\ s_j \in S_j) > 1/2.$$

But conditional on the event $(s_i \in S_i,\ s_j \in S_j)$, the two states $\omega = i$ and $\omega = j$ are exhaustive, so both inequalities cannot hold simultaneously. This is a contradiction.

Therefore, in any Bayes-Nash equilibrium, at least one player chooses N for each signal. Thus trade cannot occur in equilibrium. □

A.2 Proof of Proposition 2

Proof. Part (a). Under risk neutrality,

$$\Delta_{i,j} = \tau_{i,j}\left[2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f\right].$$

Thus, choosing $a_i = T$ upon receiving $s_i = j$ is optimal if and only if

$$2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f \geq 0,$$

whenever $\tau_{i,j} > 0$ (if $\tau_{i,j} = 0$, then $\Delta_{i,j} = 0$ and T is weakly optimal).

Under incomplete reasoning,

$$\pi_{i,j} = \alpha\mu + (1 - \alpha)\pi^*_{i,j}.$$

Substituting,

$$2(\overline{v} - \underline{v})\big[\alpha\mu + (1 - \alpha)\pi^*_{i,j} - 0.5\big] - f \geq 0,$$

which is equivalent to

$$\alpha(\mu - \pi^*_{i,j}) \geq \frac{f}{2(\overline{v} - \underline{v})} + 0.5 - \pi^*_{i,j}.$$

Hence, provided $\mu > \pi^*_{i,j}$,

$$\alpha \geq \frac{\frac{f}{2(\overline{v} - \underline{v})} + 0.5 - \pi^*_{i,j}}{\mu - \pi^*_{i,j}}.$$

Finally, the assumption $\mu \geq 0.5 + \frac{f}{2(\overline{v} - \underline{v})}$ ensures that the right-hand side is at most 1, so such an $\alpha \in [0, 1]$ exists. This proves part (a).

Part (b). Consider the strategy profile in which each player i chooses $a_i = T$ if and only if $s_i = j$.

If player j follows this strategy, then $a_j = T$ if and only if $s_j = i$. Hence, conditional on $s_i = j$, a trade occurs only when $s_j = i$, so

$$\pi^*_{i,j} = \Pr(\omega = j \mid s_i = j,\ s_j = i) = \frac{\mu(1-\mu)}{\mu(1-\mu) + (1-\mu)\mu} = \frac{1}{2}.$$

Therefore,

$$\pi_{i,j} = \alpha\mu + (1-\alpha)\frac{1}{2}.$$

Moreover,

$$\tau_{i,j} = \Pr(a_j = T \mid s_i = j) = \Pr(s_j = i \mid s_i = j) > 0.$$

Substituting into $\Delta_{i,j}$,

$$\Delta_{i,j} = \tau_{i,j}\left[2(\overline{v} - \underline{v})\alpha\,(\mu - 0.5) - f\right].$$

Thus $\Delta_{i,j} \geq 0$ whenever

$$\alpha \geq \frac{f}{2(\overline{v} - \underline{v})(\mu - 0.5)}.$$

Now consider $s_i = i$. If $a_j = T$, then necessarily $s_j = i$, so

$$\pi^*_{i,i} = \Pr(\omega = j \mid s_i = i,\ s_j = i) = \frac{(1-\mu)^2}{\mu^2 + (1-\mu)^2} < \frac{1}{2}.$$

Conditioning only on $s_i = i$ gives $\Pr(\omega = j \mid s_i = i) = 1 - \mu < 1/2$. Hence any convex combination of these beliefs is also below $1/2$, so $\pi_{i,i} < 1/2$. It follows that

$$\Delta_{i,i} < 0$$

whenever $\tau_{i,i} > 0$, and if $\tau_{i,i} = 0$ then N is weakly optimal. Thus choosing $a_i = N$ when $s_i = i$ is optimal.

Therefore, under the stated condition on α, each player optimally chooses T if and only if $s_i = j$, and this strategy profile constitutes an α-cursed equilibrium.

□

A.3 Proof of Proposition 3

Proof. Part (a). Let A denote the event that player j chooses $a_j = T$. Given player i's beliefs about player j's strategy,

$$\Pr(A \mid \omega = j) = b_{j,i}(1-\mu) + b_{j,j}\mu$$

and

$$\Pr(A \mid \omega = i) = b_{j,i}\mu + b_{j,j}(1-\mu).$$

If player i receives signal $s_i = j$, inattentive updating implies

$$\pi_{i,j} = \Pr(\omega = j \mid s_i = j, A) = \frac{\mu^\beta \left[b_{j,i}(1-\mu) + b_{j,j}\mu\right]}{\mu^\beta \left[b_{j,i}(1-\mu) + b_{j,j}\mu\right] + (1-\mu)^\beta \left[b_{j,i}\mu + b_{j,j}(1-\mu)\right]}.$$

Under risk neutrality,

$$\Delta_{i,j} = \tau_{i,j}\left[2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f\right].$$

Thus, choosing $a_i = T$ upon receiving $s_i = j$ is optimal whenever

$$\pi_{i,j} \geq 0.5 + \frac{f}{2(\overline{v} - \underline{v})}.$$

Using the expression for $\pi_{i,j}$, this condition is equivalent to

$$\frac{\pi_{i,j}}{1 - \pi_{i,j}} \geq \frac{\overline{v} - \underline{v} + f}{\overline{v} - \underline{v} - f}.$$

Substituting the odds ratio implied by inattentive updating gives

$$\left(\frac{\mu}{1-\mu}\right)^\beta \frac{b_{j,i}(1-\mu) + b_{j,j}\mu}{b_{j,i}\mu + b_{j,j}(1-\mu)} \geq \frac{\overline{v} - \underline{v} + f}{\overline{v} - \underline{v} - f}.$$

Letting $SNR \equiv \frac{\mu}{1-\mu}$ and rewriting the fraction on the left-hand side,

$$\frac{b_{j,i}(1-\mu) + b_{j,j}\mu}{b_{j,i}\mu + b_{j,j}(1-\mu)} = \frac{1 + (b_{j,j}/b_{j,i})SNR}{SNR\left[1 + (b_{j,j}/b_{j,i})(1/SNR)\right]}.$$

Hence the previous condition becomes

$$SNR^{\beta-1} \geq \frac{\overline{v} - \underline{v} + f}{\overline{v} - \underline{v} - f} \frac{1 + (b_{j,j}/b_{j,i})(1/SNR)}{1 + (b_{j,j}/b_{j,i})SNR}.$$

Since $SNR > 1$, taking logs yields

$$\beta \geq 1 + \frac{\log\left(\frac{\overline{v} - \underline{v} + f}{\overline{v} - \underline{v} - f} \frac{1 + (b_{j,j}/b_{j,i})(1/SNR)}{1 + (b_{j,j}/b_{j,i})SNR}\right)}{\log(SNR)}.$$

This proves part (a).

Part (b). Consider the candidate strategy profile in which each player i chooses $a_i = T$ if and only if $s_i = j$.

Under this strategy profile, from player i's perspective player j chooses T if and only if $s_j = i$. Thus, in part (a) we have

$$b_{j,i} = 1 \qquad \text{and} \qquad b_{j,j} = 0.$$

Substituting these values into the condition from part (a) yields

$$\beta \geq 1 + \frac{\log\left(\frac{\overline{v}-\underline{v}+f}{\overline{v}-\underline{v}-f}\right)}{\log(SNR)}.$$

Hence, under this condition, choosing $a_i = T$ upon receiving $s_i = j$ is optimal.

It remains to verify that choosing $a_i = N$ upon receiving $s_i = i$ is also optimal. Under the candidate strategy profile, if player i observes $a_j = T$, then she infers $s_j = i$. Therefore inattentive updating implies

$$\Pr(\omega = j \mid s_i = i, a_j = T) = \frac{(1-\mu)^\beta(1-\mu)}{(1-\mu)^\beta(1-\mu) + \mu^\beta \mu}.$$

Since $\mu > 1/2$ and $\beta \geq 1$, this posterior is strictly smaller than $1/2$. It follows that

$$2(\overline{v} - \underline{v})\left[\Pr(\omega = j \mid s_i = i, a_j = T) - 0.5\right] - f < 0,$$

and therefore

$$\Delta_{i,i} < 0.$$

So choosing $a_i = N$ upon receiving $s_i = i$ is optimal.

We have shown that, under the stated condition on β, each player optimally chooses T if and only if $s_i = j$. Therefore there is a Bayes-Nash equilibrium with this strategy profile.

□

A.4 Proof of Proposition 4

Proof. Part (a). Under the logit specification,

$$\sigma_{i,k} = \Pr(a_i = T \mid s_i = k) = \frac{e^{\Delta_{i,k}/\gamma}}{1 + e^{\Delta_{i,k}/\gamma}}.$$

Holding fixed player i's beliefs about player j's strategy, $\Delta_{i,j}$ is fixed. Differentiating with respect to γ gives

$$\frac{\partial \sigma_{i,j}}{\partial \gamma} = \frac{e^{\Delta_{i,j}/\gamma}}{\left(1 + e^{\Delta_{i,j}/\gamma}\right)^2}\left(-\frac{\Delta_{i,j}}{\gamma^2}\right).$$

By assumption, in the absence of cognitive frictions player i expects trade to be unprofitable upon receiving signal $s_i = j$, so $\Delta_{i,j} < 0$. Since $\gamma > 0$, it follows that

$$\frac{\partial \sigma_{i,j}}{\partial \gamma} > 0.$$

Thus the probability of choosing $a_i = T$ upon receiving signal $s_i = j$ is increasing in noise.

Part (b).

Fix any $K \in (0, 1/2)$. In a quantal response equilibrium,

$$\sigma_{i,j} = \frac{e^{\Delta_{i,j}/\gamma}}{1 + e^{\Delta_{i,j}/\gamma}}.$$

Under risk neutrality,

$$\Delta_{i,j} = \tau_{i,j} \left[2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f \right].$$

Since $\tau_{i,j} \in [0, 1]$ and $\pi_{i,j} \in [0, 1]$, it follows that

$$2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f \geq -(\overline{v} - \underline{v}) - f,$$

and therefore

$$\Delta_{i,j} \geq -(\overline{v} - \underline{v} + f).$$

Hence

$$\sigma_{i,j} \geq \frac{e^{-(\overline{v} - \underline{v} + f)/\gamma}}{1 + e^{-(\overline{v} - \underline{v} + f)/\gamma}}.$$

Now define

$$\bar{\gamma}(K) \equiv \frac{\overline{v} - \underline{v} + f}{\log\left(\frac{1-K}{K}\right)}.$$

Because $K < 1/2$, we have $\log\left(\frac{1-K}{K}\right) > 0$, so $\bar{\gamma}(K) > 0$ is well defined. If $\gamma \geq \bar{\gamma}(K)$, then

$$-\frac{\overline{v} - \underline{v} + f}{\gamma} \geq -\log\left(\frac{1-K}{K}\right) = \log\left(\frac{K}{1-K}\right).$$

Since the logistic function is increasing, this implies

$$\frac{e^{-(\overline{v} - \underline{v} + f)/\gamma}}{1 + e^{-(\overline{v} - \underline{v} + f)/\gamma}} \geq K.$$

Therefore $\sigma_{i,j} \geq K$.

Thus, for any $K \in (0, 1/2)$, if γ is sufficiently large then in any quantal response equilibrium each player trades upon receiving signal $s_i = j$ with probability at least K.

□

A.5 Proof of Proposition 5

Proof. Consider the candidate strategy profile in which each player i chooses $a_i = T$ if and only if $s_i = j$.

Under risk neutrality,

$$\Delta_{i,j} = \tau_{i,j}\left[2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f\right].$$

Thus, choosing $a_i = T$ upon receiving signal $s_i = j$ is optimal whenever

$$\pi_{i,j} \geq 0.5 + \frac{f}{2(\overline{v} - \underline{v})},$$

provided $\tau_{i,j} > 0$.

Under the candidate strategy profile, if player i observes that player j trades, then player i infers that $s_j = i$. Given player i's subjective beliefs about signal accuracies, the posterior probability that the state is j is therefore

$$\pi_{i,j} = \Pr(\omega = j \mid s_i = j,\ s_j = i) = \frac{\hat{\mu}(1 - \tilde{\mu})}{\hat{\mu}(1 - \tilde{\mu}) + (1 - \hat{\mu})\tilde{\mu}}.$$

Hence the condition for trade to be optimal is

$$\frac{\hat{\mu}(1 - \tilde{\mu})}{\hat{\mu}(1 - \tilde{\mu}) + (1 - \hat{\mu})\tilde{\mu}} \geq 0.5 + \frac{f}{2(\overline{v} - \underline{v})}.$$

Rearranging,

$$\frac{\hat{\mu}(1 - \tilde{\mu})}{(1 - \hat{\mu})\tilde{\mu}} \geq \frac{\overline{v} - \underline{v} + f}{\overline{v} - \underline{v} - f}.$$

Using the definitions

$$S\hat{N}R \equiv \frac{\hat{\mu}}{1 - \hat{\mu}} \qquad \text{and} \qquad S\tilde{N}R \equiv \frac{\tilde{\mu}}{1 - \tilde{\mu}},$$

this becomes

$$\frac{S\hat{N}R}{S\tilde{N}R} \geq \frac{\overline{v} - \underline{v} + f}{\overline{v} - \underline{v} - f}.$$

Thus, under the stated condition, choosing $a_i = T$ upon receiving $s_i = j$ is optimal.

It remains to verify that choosing $a_i = N$ upon receiving $s_i = i$ is also optimal. Under the candidate strategy profile, if player j trades then $s_j = i$, so player i's posterior belief that the state is j is

$$\pi_{i,i} = \Pr(\omega = j \mid s_i = i,\ s_j = i) = \frac{(1 - \hat{\mu})(1 - \tilde{\mu})}{(1 - \hat{\mu})(1 - \tilde{\mu}) + \hat{\mu}\tilde{\mu}}.$$

Since $\hat{\mu} > 1/2$ and $\tilde{\mu} > 1/2$, we have

$$(1 - \hat{\mu})(1 - \tilde{\mu}) < \hat{\mu}\tilde{\mu},$$

and therefore

$$\pi_{i,i} < \frac{1}{2}.$$

It follows that

$$2(\overline{v} - \underline{v})(\pi_{i,i} - 0.5) - f < 0,$$

so that

$$\Delta_{i,i} < 0$$

whenever $\tau_{i,i} > 0$, and if $\tau_{i,i} = 0$ then N is weakly optimal. Thus choosing $a_i = N$ upon receiving $s_i = i$ is optimal.

We have shown that, under the stated condition, each player optimally chooses T if and only if $s_i = j$. Therefore there is a Bayes-Nash equilibrium with this strategy profile.

□

B Calculations for the reduced-form analysis

In this appendix, we calculate a number of statistics using data from the JOINT and CURSED treatment (see Section 4).

First, we compute the probability that two trading partners have received contradicting information about the state when they agree to trade in treatment JOINT:

$$\Pr[s_i = j, s_j = i | a_i = T, a_j = T] =$$

$$= \frac{\Pr[a_i = T, a_j = T | s_i = j, s_j = i] \times \Pr[s_i = j, s_j = i]}{\sum_{s_i,s_j} \Pr[a_i = T, a_j = T | s_i, s_j] \times \Pr[s_i, s_j]} =$$

$$= \frac{\Pr[a_i = T | s_i = j] \times \Pr[a_j = T | s_j = i] \times \Pr[s_i = j, s_j = i]}{\sum_{s_i,s_j} \Pr[a_i = T | s_i] \times \Pr[a_j = T | s_j] \times \Pr[s_i, s_j]}$$

We calibrate the conditional trading probabilities using the rate at which subjects trade on different signals in JOINT:

$$\Pr[a_i = T | s_i = j] = 0.64$$

$$\Pr[a_i = T | s_i = i] = 0.14$$

To compute the probabilities of a signal profile $\Pr[s_i, s_j]$ we use the prior probabilities of each state (equal to 0.5) and the mean of the realized signal accuracy in skill-based signal treatments (equal to 0.67). This yields $\Pr[s_i = j, s_j = i | a_i = T, a_j = T] = 0.91$. Proceeding similarly using data from the CURSED treatment, we find that $\Pr[s_i = j, s_j = i | a_i = T, a_j = T] = 0.93$.

Next, we compute a player's expected gain from trading conditional on a signal suggesting the other player's asset is worth more, calibrated on data from the JOINT treatment. This is defined in expression (4):

$$\Delta_{i,j} = \{\tau_{i,j}\left[2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f\right]\}$$

where $\overline{v} = 10, \underline{v} = 2, f = 2$ and the probabilities

$$\tau_{i,j} \equiv \Pr[a_j = T | s_i = j]$$

$$\pi_{i,j} \equiv \Pr[\omega = j | a_j = T, s_i = j]$$

are computed using Bayes' rule and calibrated using state priors, the (average) signal accuracy and aggregate conditional trading probabilities. We obtain

$$\Delta_{i,j}^{JOINT} = -0.3414$$

Proceeding in a similar way for the CURSED treatment, we obtain:

$$\Delta_{i,j}^{CURSED} = -0.3178$$

C Experiment Instructions

Introduction

This survey consists of **three parts**: A, B and C, in this order.

- In part **A**, you will be asked to **forecast** how you and other participants will score in a test.
- You will take the **actual test in part C**.
- In part **B**, you will play **four games** that will be described in more detail later.

Bonus payment information

By answering all the questions in this survey, you will receive the reward for completing your submission. In addition, one out of five participants will receive a **bonus payment**. The amount of the bonus payment **depends on your choices** in the survey. At the end of the survey, the computer will randomly select one of the survey questions. If one of the Part B games is selected for payment, you will receive an amount of money between $10 and $0 as explained later. If any other question is selected for payment, you will be paid $4 if your answer to that question is right.

Next

C.1 Beliefs elicitation

PART A

In this part of the experiment, you have to **forecast** how you and other participants will **score in a test**. You will take the actual test in part C. The test consists of several puzzles. In each puzzle you will be shown a geometric design with a missing piece, with eight choices that fill in the piece. You will have to **choose the image best suited** to fill the white space in the drawing. An example is shown below. The test consists of **6 puzzles** with the same format as this one and varying difficulty.

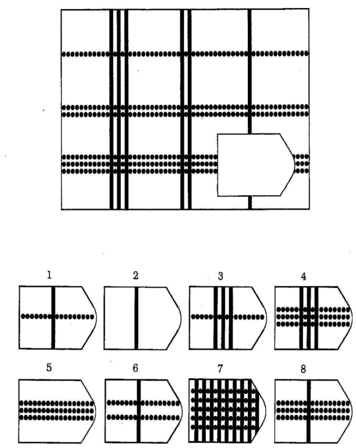

PART A

At this point, we would like you to **guess how many puzzles** you and a random other participant will answer correctly, out of the 6 puzzles in the test of Part C.

How many puzzles in the test will **you** solve correctly?

Your answer:

How many puzzles in the test will a **random other participant** solve correctly?

Your answer:

Next

C.2 Trading game instructions and comprehension check

PART B - General instructions

- In part B, you will play four games.
- Please read these instructions carefully.

Next

PART B - General instructions

- At the beginning of the game, we will randomly assign you a fictional **item**: a blue item or a red item.
- We will randomly match you with another participant (your **counterpart**) who has been assigned an item of the **opposite color**.
- That is, if you were assigned a blue item, your counterpart has been assigned a red item.
- If you were assigned a red item, your counterpart is assigned a blue item.

Previous Next

PART B - General instructions

- The computer will digitally flip a coin to determine a **winning color**: blue or red.
- That is, blue and red are equally likely to be randomly selected as the winning color.
- If the winning color is blue, then blue items will be worth \$10 and red items will be worth \$2.
- If the winning color is red, then red items will be worth \$10 and blue items will be worth \$2.

Previous Next

PART B - General instructions

- The winning color is exactly the same for both you and your counterpart, but **neither you nor your counterpart know** the winning color with certainty.

- You and your counterpart will each receive a **clue about the winning color**. Each clue is the color of a ball (either blue or red) randomly drawn from an urn:

  - The urn is filled with **100 balls**.
  - Each ball can be either blue or red.
  - If the actual winning color is red, the number of red balls in the urn will be greater than the number of blue balls.
  - If the actual winning color is blue, the number of blue balls in the urn will be greater than the number of red balls.
  - Your clue is the color of a **randomly drawn ball** from this urn.

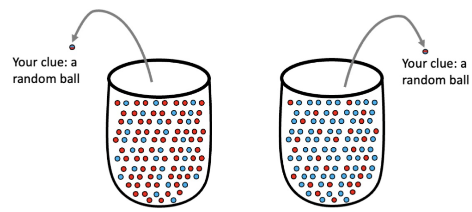

If the winning color is red   If the winning color is blue

PART B - General instructions

- Your task is to decide whether you would like to **trade your item for your counterpart's** item.

- Before making your trade decision, you will be informed whether your current item is blue or red, but we won't tell you your clue about the winning color.

- Instead, we will allow you to **plan whether to trade or not** depending on your clue.

- Here is an example of the interface that you will use to make your plan (you will make your actual choices later):

If your clue is red:	If your clue is blue:
○ I do not want to trade ○ I want to trade	○ I do not want to trade ○ I want to trade

- After you have made your plan, the computer will take your actual clue and use that along with your plan to determine your trade decision. It will do the same for your counterpart.

- If you and your counterpart **both decide to trade** (given your actual clues), then you will **exchange items**.

PART B - General instructions

- If you and your counterpart exchange items, you will each pay a $2 **fee for completing the trade**.
- In each game, you will earn a **payoff** equal to the value of **your final item** (which depends on the actual winning color) **minus the fee** (if a trade was executed).
- You will not see the outcome of the games during the experiment. Instead, we will record your and your counterpart's trading plans and determine the outcome of each game **after the experiment is over**. If one of the four games is selected for your bonus payment, your **bonus payment** will be equal to the payoff you earned in that game.

Previous Next

PART B - General instructions

Before continuing with the experiment, please answer the following comprehension questions. Before submitting your answers, you may review the Part B instructions presented in previous pages.

1) You have a blue item. The winning color is red: red items will be worth $10 and blue items will be worth $2. The fee for completed trades is $2. How many dollars will you earn in this game if both you and your counterpart **decide to trade** items?

Your answer:

----

2) You have a blue item. The winning color is red: red items will be worth $10 and blue items will be worth $2. The fee for completed trades is $2. How many dollars will you earn in this game if either you or your counterpart **decide not to trade** items?

Your answer:

----

3) If the winning color is red, then your clue is **more likely** to be red than blue.

Your answer:

Previous Next

C.3 CURSED treatment

PART B - GAME 1 Instructions

Here are some important details about the rules for **Game 1**.

About the clues

Your clue is the color of a randomly drawn ball from an urn filled with **100 balls**. In this game, there are exactly **70 balls of the winning color** and 30 balls of the other color. The composition of the urn is the **same for all participants**, and each participant's clue is an independent random draw from the same urn.

About the trading plan

In this game, your trade plans depend only on **your clue** about the winning color.

Here is an example of the interface that you will use to make your plan (you will make your actual choices in the next page):

If your clue is red:

- I do not want to trade
- I want to trade

If your clue is blue:

- I do not want to trade
- I want to trade

You have a blue item.

Value of items with winning color: $10.

Value of items with losing color: $2.

Fee for completed trades: $2.

Clue = color of a randomly drawn ball from an urn with 100 balls.

The urn composition is the same for everyone:

the urn contains 70 balls of the winning color.

Please plan your trade decision in each of the following cases:

If your clue is red:

- ○ I do not want to trade
- ○ I want to trade

If your clue is blue:

- ○ I do not want to trade
- ○ I want to trade

C.4 JOINT treatment

About the clues

Your clue is the color of a randomly drawn ball from an urn filled with **100 balls**. In this game, the **number of balls with the winning color** in your urn depends on your future **test score from part C** (the number of puzzles that you will solve correctly), as shown in the following table.

Part C test score	0	1	2	3	4	5	6
# balls with winning color	50	57	64	71	78	85	92

The higher is your future test score, the higher the number of balls of the actual winning color in the urn will be, as illustrated in the figure below. Thus, if your **test score is higher**, your clue about the winning color is **more likely to be correct**. The same applies to your counterpart. Each participant's clue is an independent random draw from that participant's urn.

Two examples where the winning color is blue

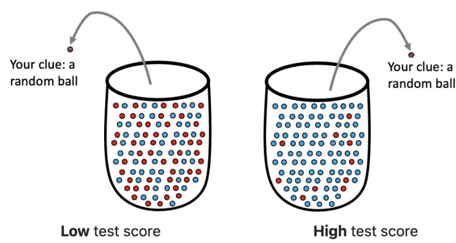

About the trading plan

Just like in Game 1, in this game your trade plans depend only on **your clue** about the winning color.

You have a red item.

Value of items with winning color: $10.

Value of items with losing color: $2.

Fee for completed trades: $2.

Clue = color of a randomly drawn ball from an urn with 100 balls.

The urn composition depends on the future test score:

Part C test score	0	1	2	3	4	5	6
# balls with winning color	50	57	64	71	78	85	92

Please plan your trade decision in each of the following cases:

If your clue is red:

- ◯ I do not want to trade
- ◯ I want to trade

If your clue is blue:

- ◯ I do not want to trade
- ◯ I want to trade

C.5 CONFIDENT treatment

About the clues

Just like in Game 2, the **number of balls with the winning color** in a participant's urn depends on that participant's future **test score from part C**.

About the trading plan

In this game, your trade plans depend on **both your and your counterpart's clues**.

Here is an example of the interface that you will use to make your plan (you will make your actual choices in the next page):

If your clue is red and your counterpart's clue is red:

- I do not want to trade
- I want to trade

If your clue is red and your counterpart's clue is blue:

- I do not want to trade
- I want to trade

If your clue is blue and your counterpart's clue is red:

- I do not want to trade
- I want to trade

If your clue is blue and your counterpart's clue is blue:

- I do not want to trade
- I want to trade

PART B - GAME 3

You have a red item.

Value of items with winning color: $10.

Value of items with losing color: $2.

Fee for completed trades: $2.

Clue = color of a randomly drawn ball from an urn with 100 balls.

The urn composition depends on the future test score:

Part C test score	0	1	2	3	4	5	6
# balls with winning color	50	57	64	71	78	85	92

Please plan your trade decision in each of the following cases:

If your clue is red and your counterpart's clue is red:

- I do not want to trade
- I want to trade

If your clue is red and your counterpart's clue is blue:

- I do not want to trade
- I want to trade

If your clue is blue and your counterpart's clue is red:

- I do not want to trade
- I want to trade

If your clue is blue and your counterpart's clue is blue:

- I do not want to trade
- I want to trade

C.6 CLEAN treatment

PART B - GAME 4 Instructions

Here are some important details about the rules for **Game 4**.

About the clues

Just like in Game 1, the urn is filled with **70 balls of the winning color** and 30 balls of the other color.

About the trading plan

Just like in Game 3, your trade plans depend on **both your and your counterpart's clues**.

When you are ready to play the game, click "Next".

Next

PART B - GAME 4

You have a red item.

Value of items with winning color: $10.

Value of items with losing color: $2.

Fee for completed trades: $2.

Clue = color of a randomly drawn ball from an urn with 100 balls.

The urn composition is the same for everyone:

the urn contains 70 balls of the winning color.

Please plan your trade decision in each of the following cases:

If your clue is red and your counterpart's clue is red:

- ○ I do not want to trade
- ○ I want to trade

If your clue is red and your counterpart's clue is blue:

- ○ I do not want to trade
- ○ I want to trade

If your clue is blue and your counterpart's clue is red:

- ○ I do not want to trade
- ○ I want to trade

If your clue is blue and your counterpart's clue is blue:

- ○ I do not want to trade
- ○ I want to trade

Next

C.7 Raven's test

PART C

Starting from the next page, you will see **6 puzzles** (one for each page). For each puzzle, choose from the 8 options displayed on the left of the screen the image best suited to fill the white space in the drawing.

Next

PART C - PUZZLE 1

Choose from the 8 options displayed at the bottom of the drawing the image best suited to fill the white space in the drawing.

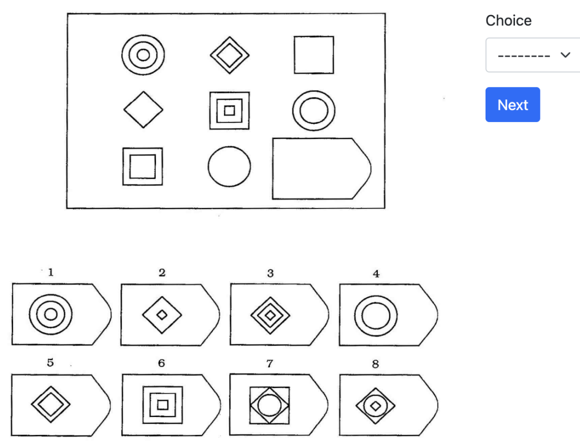

D Earlier Working Paper

The current paper replaces and subsumes our earlier working paper, "Why Do People Violate No Trade Theorems? A Diagnostic Test." After receiving comments on that paper we decided to replace it with a new experiment that removes some shortcomings and ambiguities of that earlier design. This appendix summarizes the design and main findings of our earlier experiment and highlights the key differences relative to the current study.

Our earlier experiment used the same basic trading game as our current study: two players receive private signals about an underlying state, and make a simultaneous decision of whether to swap assets. As in the current design, we implemented the game using the strategy method, eliciting subjects' willingness to trade contingent on each possible signal realization. The main differences between the two designs concern the signal-generating task and the treatments. In the previous experiment, we used a different task to generate skill-based signals. Subjects were briefly shown 400 colored dots, with a slight majority of one color depending on the state, and must estimate the state from this display. The limited viewing time prevented exact counting and induced a noisy estimate. Like in the current experiment, the treatment structure was designed to isolate different mechanisms that could generate trade, but using different manipulations. First, in some sessions, we eliminated scope for relative overconfidence by replacing subjects' own estimates with estimates generated by third parties, which were then used to implement the subjects' strategies. Second, some subjects were informed of their counterpart's signal-contingent strategy prior to making their own decision, with the aim of limiting the scope for incomplete reasoning.

The results are broadly similar to those reported here. Subjects trade at high rates in baseline conditions, with more than 70% of subjects choosing to trade. Treatment comparisons and structural estimates suggest that (what we call here) incomplete and imprecise reasoning along with relative overconfidence play important roles in generating trade, and that overconfidence appears to manifest as dismissiveness regarding the accuracy of counterparties' estimates. Our conclusions about overconfidence in that earlier paper however were derived purely from our finding of subjects overweighting their own signals and (unlike in the current experiment) we did not collect belief data to allow us to test whether this was driven by overconfidence or own-signal salience.

The current experiment builds on this earlier design but introduces several important improvements. First, a key limitation of the signal-generating task in our previous experiment is that the experimenter does not directly observe the signal received by participants. Instead, one only observes subjects' estimates and can construct a measure of realized accuracy ex post. The estimate itself, however, may reflect not only the subject's information but also strategic reporting considerations. In contrast, in the current design we measure subjects' skills separately and then generate signals based on these skills, allowing us to observe both the signal and its accuracy directly.

Second, this limitation is compounded by the fact that (as discussed above) the previous experiment did not elicit subjects' beliefs about signal accuracy. As a result, overconfidence could only be inferred indirectly from behavior. In the current experiment, we elicit beliefs about accuracy, which allows us to directly measure overconfidence and relate it to trading behavior.

Third, to eliminate scope for relative overconfidence, the previous experiment relied on treatments in which signals were based on third-party estimates. While this removes the link between a subject's own skill and their signal, it introduces a new potential confound: subjects' perceptions of the accuracy of third-party estimates. In the current design, we instead use truly exogenous signals, which eliminates scope for overconfidence in a more direct and transparent way and allows us to precisely control signal accuracy.

Fourth, to address failures of strategic reasoning, the previous experiment included a treatment in which subjects were informed of their counterpart's signal-contingent strategy. The goal was to reduce the scope for cursed reasoning. However, recent evidence suggests that such interventions may not fully resolve the underlying inference problem. In the current experiment, we address this issue more directly by making signals public in one treatment, thereby removing the need for subjects to infer information from others' actions. This makes the assumption that the treatment eliminates scope for incomplete reasoning more credible.

Fifth, adding the public-signal treatment in the current experiment provides an additional advantage by allowing us to observe subjects' choices in two critical cases: when their own signal suggests trade but their counterpart's signal does not, and vice versa. These comparisons are essential for identifying how subjects weight their own information relative to others'. In the previous design, even when subjects were informed of their counterpart's signal-contingent strategy, we could infer behavior conditional on at most one realization of the counterpart's signal (namely, the one on which the counterpart chose to trade). As a result, the earlier experiment could not directly identify the role of selective attention to one's own signal.

Finally, the previous experiment varied the level of transaction fees across rounds, while the current experiment uses a single fee throughout. This difference is not central to the comparison: variation in transaction costs could, in principle, be reintroduced in the current design without affecting the main identification strategy.