Why Do People Violate No-Trade Theorems?
A Diagnostic Test

Jacopo Magnani   Ryan Oprea*

December 15, 2016

Abstract

Evidence from both the lab and field suggest that people frequently trade on the basis of private information in violation of no-trade theorems. Why? We report an experiment designed to distinguish three prominent explanations: relative overconfidence about private information (as emphasized in the behavioral finance literature), limited strategic sophistication (as emphasized recently in the behavioral economics literature) and noisy best response to weak incentives (as emphasized more generally in the experimental economics literature). Our experimental design takes a diagnostic approach, stripping these potential explanations away one-by-one over a series of treatments. We find that when all of these channels are available as explanations, subjects agree to trade in excess of 70% of the time. When we remove them all, violations of no-trade theorems disappear entirely. Treatment level analysis and structural estimation suggests that relative overconfidence and strategic sophistication both have large and roughly equal effects on motivating trade and that noisy behavior can have a large additional effect when the costs of violating no-trade theorems is small.

Keywords: TBA

JEL codes: C9, D03 G1

1 Introduction

For over thirty years economists have known that speculative trade based on private information should not occur in markets. As "no-trade theorems" (e.g. Milgrom and Stokey (1982), Tirole

*Magnani: Purdue University, West Lafayette, IN, 47906, jacopo.magnani@gmail.com; Oprea: Economics Department, University of California, Santa Barbara, Santa Barbara, CA, 95064, roprea@gmail.com.

(1982)) point out, a trader whose private information suggests an asset is worth buying can only execute a trade if there exists a seller with private information suggesting the opposite –the existence of such a willing seller will counteract the potential buyer's private information, causing her to choose not to trade in equilibrium.[1] Nonetheless, there is ample evidence from both the field (Odean (1999))) and the lab (Angrisani et al. (2008), Carrillo and Palfrey (2011)) that speculative efforts to trade on the basis of private information are a common feature of human economic behavior. Observers have noted that the volume and types of trade observed in the field cannot be (entirely) explained by insurance and liquidity motives[2] and laboratory experiments that carefully control these alternative motives show robust evidence of trade (Angrisani et al. (2008), Carrillo and Palfrey (2011)). The evidence thus suggests that people routinely violate no-trade theorems.

Why do people engage in trade in settings where the no-trade theorems apply? A number of explanations are available. Researchers in behavioral finance (e.g. Odean (1998), Daniel et al. (2001), Scheinkman and Xiong (2003); see Daniel and Hirshleifer (2015) for an overview) have long emphasized the potential role of relative overconfidence: believing one's information is "better" than that of trading partners breaks the common priors assumption underlying no trade theorems, generating scope for trade. Researchers in behavioral economics, on the other hand, have recently emphasized limitations in strategic reasoning as an explanation. In *cursed belief* models buyers do not fully appreciate the degree to which willingness to trade by sellers is conditioned on sellers' information and vice versa (Eyster et al. (2015), Carrillo and Palfrey (2011)).[3] Finally, noisy trading behavior (due to bounded rationality or idiosyncratic preferences for trade) can have both direct and equilibrium effects that generate willingness to trade: noisy equilbrium models such as the Quantal Response Equilibrium (QRE, McKelvey and Palfrey (1995), Goeree et al. (2016)) model can break the linkage between information and willingness to trade, again inspiring violations of no-trade theorems.

We report an experiment designed to evaluate these three channels of explanation and discover

[1]This is true under the assumption there are no risk-sharing motivations for trade.

[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."

[3]Similarly, limited depth of reasoning models (Nagel (1995), Stahl and Wilson (1995), Camerer et al. (2004)) such as Level-k models and Cognitive Hierarchy models can generate scope for trade. In this paper we focus on cursed reasoning models but emphasize that many of our results can also be rationalized using these closely related alternatives.

why people trade based on private information. Subjects are each assigned an Arrow security and, based on private information about the state, must decide whether they are willing to pay a transaction fee to swap securities with a counterpart who holds a security that may have a higher or lower value, depending on the state. The diagnostic method we pursue in our experimental design is to strip away each of the above explanations as plausible mechanisms and examine whether violations of no-trade theorems persist (our design rules out risk-preference-based explanations altogether). In some of our treatments subjects form their beliefs about the state based on an estimation task that depends to some degree on the subject's skill, generating scope for relative overconfidence, but in further treatments we eliminate relative overconfidence as a source for trade by instead providing both subjects with estimates made by third parties.[4] Likewise in some of our treatments subjects must form conjectures about the sophistication of their counterparts' trading strategies, generating scope for cursed beliefs while in others we eliminate this possibility by showing subjects the strategies of their potential trading partners directly before they make their own trading decisions. Finally, we study a trading game specifically designed to make the costs of violating no trade theorems completely exogenous, allowing us to directly control and vary the plausibility of noisy trade models like QRE.

We find large treatment effects in our data. In the baseline condition, in which costs of violating no trade theorems are modest and relative overconfidence and cursed reasoning are possible, most subjects – over 70% – agree to trade in violation of no-trade theorems. When we remove relative overconfidence and cursed reasoning as possibilities and parameterize games so that the costs of deviating from best response are non-negligible, people stop violating no-trade theorems *entirely*: we observe zero trade in such settings exactly as predicted by no-trade theorems. Because we can turn these explanatory channels on and off in our design, we can pinpoint exactly how various channels influence behavior. Adding scope for relative overconfidence or cursed reasoning each add about 30 percentage points to the rate at which agents agree to trade and reducing the cost of violations to low levels in a way that invites noisy trading behavior adds another 30-40 percentage points.

Structural estimates reveal that subjects employ extremely strong cursed reasoning, suggesting that they believe that their counterparts are over 80% likely to neglect private information in making trade decisions. Estimates also reveal that subjects are highly relatively overconfident, but that this relative overconfidence takes a special form. While estimated parameters suggest

[4] Moore and Healy (2008) refer to relative overconfidence as *overplacement*, and emphasize that it is conceptually different from *overestimation* of one's actual perfomance. Indeed, one of our findings is that subjects show evidence of suffering from overplacement in our data but not overestimation.

subjects form quite accurate beliefs about their own average ability to infer the state, they suggest that subjects treat their counterparts' inferences as little better than white noise. Eyster et al. (2015) call this type of relative overconfidence *dismissiveness* and note that it acts as a substitute rather than a complement to cursed reasoning. The effect of this in our data is that subjects do not violate no-trade theorems with much greater frequency when both relative overconfidence and cursed reasoning are allowed by the design than when either is allowed alone.

Our research builds on a small prior experimental literature studying no-trade theorems. Two pioneering papers are particularly relevant: Angrisani et al. (2008) and Carrillo and Palfrey (2011) both study bilateral trading 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 in that we designed our experiment to identify sources of violations of no-trade theorems by (i) studying a game designed to make this sort of inference particularly clear[7] and (ii) by studying a set of treatments engineered specifically to remove, one-by-one, hypothesized channels by which such

[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 Ngangoue and Weizsacker (2015) 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.

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

Understanding which channels drive trade is important both for theoretical reasons and policy reasons. On the theoretical side, cursed reasoning and relative overconfidence each can have dramatically different implications for large-scale markets, as Eyster et al. (2015) point out. For instance, Eyster et al. (2015) study competitive models in which excessive speculative trade dies out as markets grow large if it is driven by overconfidence, but can survive in very large markets when driven by cursed reasoning. Our finding that speculative trade is directly linked to cursed reasoning and not only overconfidence suggests that direct evidence on no-trade violations collected in the laboratory is likely to scale in larger markets. Likewise, Eyster et al. (2015) point out that cursed reasoning can be a propogation mechanism allowing relative overconfidence to survive and further increase trade in large market if it takes the form of (absolute) overconfidence but cannot serve the same function if it takes the form of very extreme dismissiveness. Our finding that trade linked to relative overconfidence is driven by extreme dismissiveness (rather than overconfidence) suggests that relative overconfidence is unlikely to compound the excess trade generated by cursed reasoning in large markets.

On the policy side, violations of no-trade theorems can have sizable negative effects on traders' wealth (e.g. Barber and Odean (2000)) and distortionary effects on market outcomes (e.g. Daniel and Hirshleifer (2015), Scheinkman and Xiong (2003), Eyster et al. (2015)). Understanding the mechanisms that lead to these violations is important for designing appropriate interventions, such as nudges and financial education. If relative overconfidence is a primary driver of trade, then influencing traders to process information in a more objective manner could in principle lead them to avoid harmful trades, but this sort of policy could be ineffective if the underlying issue is cursed reasoning. Our results suggest that interventions that tackle relative overconfidence without tackling limited strategic thinking (or vice versa) are unlikely to be effective since both operate as strong motivations for trade. Our results also suggest that the form of relative overconfidence that generates trade in our data is, like cursed reasoning, due to underestimation of trading partners. Together our results therefore suggest that interventions aimed at raising estimation of trading partners rather than reducing hubris about one's own abilities may be more effective in reducing irrational trade.

The remainder of the paper is organized as follows. In Section 2 we describe the model we study in the lab, derive a no-trade theorem for the model and show how relative overconfidence, cursed reasoning and noisy trading behavior can generate equilibrium violations of no-trade theorems. In

Section 3 we discuss our experimental design. In Section 4 we provide reduced form evidence from our experiment describing our treatment effects, and structural analysis that generates parameter estimates that aid in interpreting our results. Finally, we close the paper with a discussion of our findings in Section 5. Proofs are provided in Appendix A and instructions to subjects in Online Appendix B.

2 Theory and Hypotheses

In Section 2.1 we describe the model we implement in our experiment and prove a no-trade theorem for this environment. In Section 2.2 we discuss three potential behavioral causes for violations of the no-trade theorem and state a pair 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 $\alpha_i \in \{\alpha_1, \alpha_2\}$ whose value depends on the state $\omega_i \in \{\omega_1, \omega_2\}$. Agent i's security, α_i, pays $v = \overline{v}$ if the state is ω_i and $v = \underline{v}$ if the state is $\omega_j, j \neq i$, 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 the two signals are independent conditional on the state. We define $\mu_i \in (0.5, 1)$ as the likelihood that the signal player i receives is correct:

$$\mu_i \equiv Prob(s_i = k|\omega_k) > 0.5 \text{ for } i \in \{1, 2\}, k \in \{1, 2\}$$

In order to match the setting of the experiment (see below), we assume that μ_i is unkown to any agent but is drawn from a known distribution, common to all agents. In particular we assume that each μ_i is independently and identically distributed. We denote by $\overline{\mu}$ the mean likelihood. (Making the more conventional assumption that μ_i is known and common across subjects changes nothing in what follows.)

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 (agent i's asset is then α_j) 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 $\Delta_{i,k}$:

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

where $\tau_{i,k} \equiv Prob(a_j = T | s_i = k)$ is player i's conditional belief that her counterpart will trade (i.e. that $a_j = T$) and $\pi_{i,k} \equiv Prob(\omega_j | s_i = k, a_j = T)$ is the probability she attaches to the state being ω_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 in other no-trade theorems, the basic idea is that, given the common value nature of the assets, both player cannot 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.*

2.2 Three Sources of No-Trade Theorem Violations

What deviations from standard neo-classical assumptions might cause agents to trade in an environment like this? We identify three candidate explanations from the literature that we designed our experiment to assess and that we formalize for our game below. First, as the behavioral finance literature has emphasized (Odean (1998), Daniel et al. (2001), Scheinkman and Xiong (2003)), an agent can justify trading if she believes her signal is more accurate, on average, than her counterpart's signal; we will call this channel *relative overconfidence*. As Eyster et al. (2015) point out this can take two forms, each of which can generate trade: *overconfidence* (in which an agent believes her signal is *more* accurate than it actually is) and *dismissiveness* (in which an agent believes her counterpart's signal is *less* accurate than it actually is). Second, as the behavioral economics literature has emphasized, an agent can justify trading if she believes her counterpart is strategically unsophisticated, failing to fully condition trading decisions on signal. This idea can be formalized using limited depth of reasoning models such as the Level-k model (Nagel (1995), Stahl and Wilson (1995),Camerer et al. (2004)), but we follow recent theoretical and experimental work on no-trade violations (Eyster et al. (2015), Carrillo and Palfrey (2011)) and formalize it using the closely related (in this class of games at least) cursed equilbrium model of Eyster and Rabin (2005). We will therefore call this channel *cursed reasoning.*

Finally, as the experimental economics literature has emphasized more broadly (see for example Camerer et al. (2016)), agents often make mistakes (noisily deviate from best response) in games and do so especially often when the costs of mistakes are small (i.e. due to bounded rationality or small idiosyncratic perturbations to preferences) and these mistakes can generate equilibrium effects that can, in turn, generate trade; we will call this channel *quantal response equilibrium* or *QRE*, following the most commmon method for formalizing the idea in experimental economics (McKelvey and Palfrey (1995), Goeree et al. (2016)).

We formalize these three candidates explanations below and show that they can all generate trade in our game. We also, in 2.3, consider how the underlying behavioral tendencies behind these explanations can interact to reinforce one another. In the discussion below we simplify the argument and notation by assuming players are risk-neutral. Under risk-neutrality, we can simplify the expression for player i's expected gain from trading as follows:

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

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

2.2.1 Relative Overconfidence

First, we formalize *relative overconfidence* by assuming that overconfident agent i believes the likelihood of his own signal being correct is on average higher than the likelihood of player j's signal:

$$\tilde{\mu} \equiv E^i \mu_i > E^i \mu_j \equiv \hat{\mu}$$

where E^i denotes player i's expectations.

If a player i is sufficiently relatively overconfident, trading upon receiving signal $s_i = j$ becomes an optimal strategy. 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:

$$SNR_i \equiv \frac{\tilde{\mu}}{1 - \tilde{\mu}}$$

and

$$SNR_j \equiv \frac{\hat{\mu}}{1 - \hat{\mu}}$$

We can then prove the following Proposition:

Proposition 2. *If* $\frac{SNR_i}{SNR_j} \geq \frac{\overline{v} - \underline{v} + f}{\overline{v} - \underline{v} - f}$ *it is optimal for player* i *to choose* $a_i = T$ *upon receiving signal* $s_i = j$.

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 $\hat{\mu} = 0.65$, $\tilde{\mu} = 0.82$, players adopt trading strategies for all the fee levels and asset payoffs considered in this experiment. Notice that this result does not depend on whether relative overconfidence is driven by dismissiveness (a belief that a counterpart's signal is less likely to be correct than it actually is) or overconfidence (a belief that a player's own signal is more likely to be correct than it actually is). See Eyster et al. (2015) for more on the distinction between these types of relative overconfidence.

2.2.2 Cursed Reasoning

Second, we formalize *cursed reasoning* using the notion of cursed equilibrium (Eyster and Rabin (2005)), in which each agent i believes that with probability $\chi \in [0, 1]$ the action of her counterpart is *independent* of the counterpart's signal and with probability $1 - \chi$ the counterpart trades according to her signal as in a BNE.

If player i is sufficiently cursed, trading upon receiving signal $s_i = j$ becomes an optimal strategy. When player i believes counterpart j's action is partly independent of j's signal, player i perceives a larger gain from trade than in a BNE. If the perceived gain is larger than the trading fee then player i expects to profit from a trade. To formalize this argument, we prove the following Proposition:

Proposition 3. *If* $\chi \geq \frac{f(\overline{\mu}-\overline{\mu}^2)}{2(\overline{v}-\underline{v})(\overline{\mu}-0.5)[\overline{\mu}-E(\mu_i{}^2)]}$ *there is a cursed equilibrium where player* i *chooses* $a_i = T$ *upon receiving signal* $s_i = j$*. If there is no uncertainty in the signal likelihood, the condition for a cursed equilibrium simplifies to* $\chi \geq \frac{f}{2(\overline{v}-\underline{v})(\overline{\mu}-0.5)}$

For example, suppose μ_i can take two value, 0.7 or 0.8, with equal probability. If $\chi = 0.85$, trading is a cursed equilibrium for all the fee levels and asset payoffs considered in this experiment (this is also true if there is no uncertainty over signal accuracies and μ_i is identically equal to 0.75).

2.2.3 Quantal Response

Finally, we formalize the equilibrium effects of noisy decision making on trade using the logit specification of the quantal response equilibrium (QRE) model (McKelvey and Palfrey (1995), Goeree et al. (2016)), in which the probability player i chooses to trade conditional on $s_i = k$ is:

$$\sigma_{i,k} \equiv Prob(a_i = T|s_i = k) = \frac{e^{\lambda\Delta_{i,k}}}{1 + e^{\lambda\Delta_{i,k}}}. \tag{2}$$

Here $\lambda > 0$ parameterizes the responsiveness of a player's choice to the size of the expected gains from trading. When $\lambda = 0$ behavior is completely unresponsive to expected gains and the player randomly chooses whether or not to trade: thus the conditional probability of trading is 0.5 for both signals ($\sigma_{i,k} = 0.5$ for all k). When behavior is fully responsive to expected payoffs, in the limit as $\lambda \to \infty$, we obtain the no-trade equilibrium: trading yields a loss on average and thus the trading probabilities are set to zero ($\sigma_{i,k} = 0$ for all k).

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, $\pi_{i,k}$, and also affects the conditional probability that that counterpart will agree to trade, $\tau_{i,k}$ (details on how QRE affects these parameters are provided in Section 4.3). This is because when player j's decision making is noisy, she is more likely to trade regardless of her signal, easing adverse selection in the game and reducing agent i's losses from trading. This, in turn, encourages more trade from i, generating an amplified equilibrium effect. For any finite $\lambda > 0$, trade can occur under our parameters in a quantal response equilibrium (and at an increasing rate the smaller is f).

2.3 Interactions

These channels for trade can interact in important ways and sometimes reinforce one another. For instance, when we add relative overconfidence or cursed reasoning into a quantal response model, the probability player i chooses to trade conditional on signal $s_i = j$, $\sigma_{i,j}$, increases relative to a pure QRE model. This is simply because both relative overconfidence and cursed reasoning increase the posterior belief that the true state is j, $\pi_{i,j}$, thus increasing the perceived gain from trading upon receiving signal j, $\Delta_{i,j}$.

Less obviously, the way cursed reasoning and relative overconfidence interact depends on the nature of relative overconfidence. As Eyster et al. (2015) note, cursed beliefs and overconfidence work as complements, while cursed beliefs and dismissiveness act as substitutes. To illustrate, it is convenient to ignore QRE noise and uncertainty about the signal accuracies, although the main intuition holds in a more complicated framework including these elements. In this setting, the posterior belief of a player with cursed beliefs and relative overconfidence is:

$$\pi_{i,j} = \chi\tilde{\mu} + (1-\chi)\frac{\tilde{\mu}(1-\hat{\mu})}{\tilde{\mu}(1-\hat{\mu}) + (1-\tilde{\mu})\hat{\mu}} \tag{3}$$

where $\tilde{\mu}$ and $\hat{\mu}$ are the accuracies of a player's own signal and of her counterpart's signal, as perceived by a relatively overconfident player. We can prove the following Proposition:

Proposition 4. *Cursed reasoning and overconfidence are complements in the sense that:* $\frac{\partial^2 \pi_{i,j}}{\partial\chi\partial\tilde{\mu}} \geq 0$. *Cursed reasoning and dismissiveness are substitutes in the sense that:* $-\frac{\partial^2 \pi_{i,j}}{\partial\chi\partial\hat{\mu}} \leq 0$.

The Proposition shows that the marginal effect of an increase in the perceived accuracy of a player's own signal increases with the cursedness of her beliefs and, moreover, the marginal effect of an increase in the degree of cursedness increases with the perceived accuracy of a player's own signal. On the other hand, the marginal effect of a reduction in the accuracy attributed to a counterpart's signal decreases with the cursedness of beliefs and the marginal effect of an increase in the degree of cursedness decreases with this sort of dismissiveness.

In the extreme case in which a player is *fully dismissive*, perceiving his counterpart's signal as completely uninformative, i.e. $\hat{\mu} = 0.5$, the posterior belief collapses to $\pi_{i,j} = \chi\tilde{\mu} + (1-\chi)\tilde{\mu} = \tilde{\mu}$ and cursed thinking has no additional effect above and beyond the effect of dismissive relative overconfidence. Moreover in this case a player's posterior belief is equivalent to the belief of a fully cursed player with $\chi = 1$.

3 Design

In Section 3.1 we describe our strategy for implementing the model described in Section 2.1 in the laboratory. In Section 3.2 we discuss our treatment design. Finally in Section 3.3 we give details

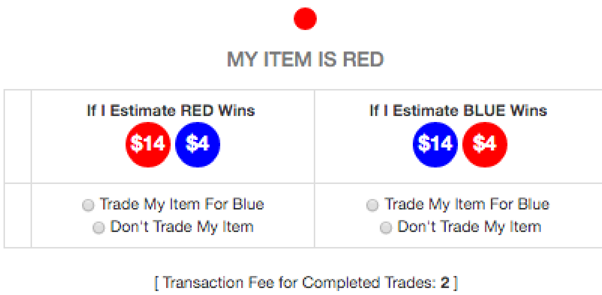

Figure 1: Screenshot for P1 types.

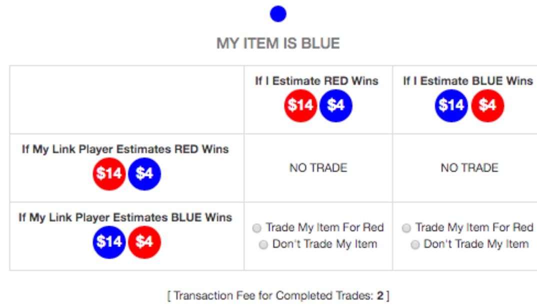

Figure 2: Screenshot for P2 types.

on how we conducted the experiment.

3.1 Overview of Experiment

The experiment is divided into several periods, each an independent run of the game described in Section 2. At the beginning of each period, the computer randomly draws a state $\omega \in \{red, blue\}$ (each occurring with equal likelihood) and assigns each subject an Arrow security that is either of a "red" or "blue" type. A security's value to subjects is $\overline{v}$ =\$14 if it is of the same color as the realized state, and $\underline{v}$ =\$4 otherwise. We do not inform subjects of the state directly but instead give them each a noisy signal as described below. Subjects are each randomly matched with a "link player," who is endowed with a security of the opposite color and given an opportunity to exchange securities with this player. If both the subject and her link player agree to trade, they will exchange securities and they will each pay a transaction fee $f \in \{0, 1, 2, 3, 4\}$ (which varies from period to period) to the experimenter.

We implement the game using two special features that allow us (using the treatments described in Section 3.2) to diagnostically vary the operation of the explanations for trading behavior considered in Section 2.2.

First the accuracy of the signal we provide to subjects depends to some degree on the subject's skill, generating scope for overconfidence. Specifically, we show subjects 400 colored dots on their screen for 10 seconds: if the state chosen by the computer is *red*, subjects are shown 205 red dots and 195 blue dots which are arrayed in a random order on the screen (and vice versa if the state is *blue*).[8] If a subject can determine which color dot is more prevelant on the screen, the subject can

[8]Caplin and Dean (2014) use a similar task to study information acquisition under rational inattention in an individual choice task. Their setup differs in minor ways from ours (number of dots on the screen, time limits etc.) but they find (as we do) that subjects behave as if they observe a noisy signal of the true state.

determine the state. We do not, however, give subjects enough time to count the dots and they therefore can only noisily estimate the state (in the data, subjects on average estimate the state with 65% accuracy). We thus give subjects noisy signals, but signals whose accuracy may vary from subject to subject based on their estimation ability, creating scope for overconfidence.[9]

Second, we implement the experiment, each period, in the reverse of the natural order, with the trading decision occurring *before* the subject sees her signal (this is necessary for one of the main treatment variations discussed in Section 3.2). Specifically, we divide the experiment into a "Trading Stage" followed by an "Estimation Stage." In the Estimation Stage subjects see their signal of dots (as described above), are each asked to make an incentivized estimate $e_i \in \{red, blue\}$ of which color dot is more prevelant and thus which state has occured (subjects earn $14 if their estimate is correct, $4 otherwise). Prior to the Estimation Stage, in the Trading Stage, subjects decide whether they are willing to trade as a function of the estimate they will later make in the Estimation Stage. That is, subject i chooses $a_i(e_i)$, a mapping $a_i : \{red, blue\} \mapsto \{T, N\}$ from a set of (future) estimates to a set of trading decisions (trade, T, or no trade, N). Subjects make this decision using a display like the one shown in Figure 1. At the top of the display, subjects are told their item color and at the bottom the period's transaction fee, f. Between is a table with two columns, each linked to a possible estimate, e_i, the subject may make in the Estimation Stage. At the period's end, after the Estimation Stage, the computer implements the decision made in the left hand column if the subject estimates $e_i = red$ and the decision from the right hand column if the subject estimates $e_i = blue$. Subjects thus make a decision for each contingency, resulting in an empirical decision function $a_i(e_i)$.

After both stages are over, the computer implements a trade (and levies a transaction fee on each party) only if both i and her link player $-i$ choose decision functions and make estimates generating outcomes $a_i(e_i) = T$ and $a_j(e_j) = T$. Subjects are given no feedback on the outcomes of trades, the realizations of states or the accuracy of estimates until all of the periods are completed and the session is finished. At the end of the experiment, subjects are paid for one randomly selected period (out of the 5 total periods) and for one randomly selected stage of the period (Estimation or Trading); they are thus paid either for a Estimation or Trading decision but not both.

[9]Importantly, as we show below, subjects are given no feedback on their accuracies and do not make decisions that suggest that they are aware of their accuracies relative to other subjects.

3.2 Treatment Design

We designed the experiment to systematically vary the operation of potential explanations for violations of the no trade theorem described in Section 2.2 one by one, over a series of treatments.

First, we remove cursed reasoning as a candidate explanation for half of our players (called "P2 types") by informing these subjects of exactly how their potential trade partner's actions depend on beliefs. Specifically, while we assigned half of subjects in each session to be "P1 types" who make their trading decisions simultaneously with their link players using the interface pictured in Figure 1, we assigned the other half of subjects to be "P2 types" who make their trading decisions only after their link players (who are always P1 types) do. These P2 types are shown their link player's decision function while making their own decisions, using an interface like the one shown in Figure 2. P2 subjects can condition trading decisions on their own e_i choices (shown as the rows of the table), but they can also sees whether their link player, j's willingness to trade is dependent on j's estimate e_j. In the example in Figure 2, the link player, j, is not willing to trade if her estimate e_j is *red* (the corresponding cells are filled out with the words "NO TRADE") but is willing to trade if her estimate is $e_j = blue$. It is thus made clear to the P2 subject that her trading decision is only implemented if $-i$ estimates that the state is *blue*. (If link player j's trading decisions are independent of her estimate, e_j, the two rows of the table are merged on subject i's screen.) Thus, while P1 types must reason about their counterparts' strategies, P2 types do not, removing the possibility that cursed reasoning drives trade for these subjects.[10]

(In order to prevent P2 subjects' access to their link players' strategies from influencing their link players' choices, we employed a special matching protocol. Pairs of P1 types are matched with one other, each serving as one another's link players. Each P2 subject also has one of these P1 types as her link player but does not serve as her link player's (or any other subject's) link player. Thus P2 types are influenced by but have no influence over the earnings of other players, preventing P2 players' access to strategy information from influencing P1 players' decisions. Subjects were made fully aware of this matching protocol and the way their decisions influenced the outcomes of others.)

Second, we remove relative overconfidence as a candidate explanation in half of our sessions – called "No Relative Overconfidence" or "NO" sessions – by de-linking subjects' own estimates of the state from their trading decisions. In our "Relative Overconfidence" or "O" sessions, subject

[10] This treatment variation also removes closely related types of reasoning such as Level-k thinking (Nagel (1995), Stahl and Wilson (1995), Camerer et al. (2004)).

i's trading decision is $a_i(e_i)$, meaning her trading decision is governed by her own estimate of the state, as described in Section 3.1. In "NO" sessions, by contrast, subject i's trading decision is $a_i(e_k)$, where k is a randomly selected participant (called subject i's "guesser") from the session (the subject also knows that the same procedure is used for her link player's trading decision). While in "O" sessions a subject may trade because she believes the accuracy of her estimate is higher than the estimate of her link player (relative overconfidence), the same is not true in the "NO" treatment where the expected accuracy of the estimate determining each player's trade decision is exactly the same. Thus in the NO treatments, relative overconfidence is eliminated as a possible cause of trade.

These two treatment variables are varied within session (P1 vs. P2 type subjects) and between session (O vs. NO type sessions) giving us four treatments as summarized in Table 1. P2 subjects in the NO condition are subject to neither cursed reasoning nor relative overconfidence and we call this the CLEAN treatment. P1 subjects in the NO condition are subject to cursed reasoning but not relative overconfidence while P2 subjects in the O treatment are subject to relative overconfidence but not cursed reasoning – we call these the CURSED and CONFIDENT treatments, respectively. Finally, P1 subjects in the O treatment are potentially subject to both cursed reasoning and relative overconfidence and we therefore call this treatment JOINT.

Third, we systematically vary the viability of noisy decision making (QRE) generating trade by directly varying the expected cost of trading. This cost is simply the transaction fee, f and we vary this cost from period-to-period between \$0 and \$4. At a fee of zero, trading is in fact weakly rational, at fees of 1 or 2 moderately irrational and at fees of 3 or 4 subjects must be willing to lose a considerable fraction of earnings relative to best response in order to execute a trade. Under the hypothesis that noisy decision and its equilibrium effects are major drivers of violations of no-trade theorems, we expect no-trade predictions to be obeyed with greater likelihood the higher these fees grow in each treatment.

Finally, it is worth emphasizing that our model and therefore our experimental design deliberately rules out *preferences* as an explanation in all of our treatments. Unlike standard trading games, where players can remove uncertainty and the risk of incurring losses by substituting cash for the 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.

Subject Type	Session Type O	NO
P1	JOINT	CURSED
P2	CONFIDENT	CLEAN

Table 1: Treatment design

3.3 Implementation

We ran the experiment with custom software programmed in Javascript in August and September of 2016 at the EBEL laboratory at the University of California, Santa Barbara using mostly undergraduate subjects recruited from across majors via ORSEE (Greiner (2015)). We seated subjects at visually isolated terminals and read them instructions (reproduced in Online Appendix B) aloud. We then ran subjects through three practice periods (without pay or feedback) and, after reading the remainder of instructions, administered a quiz to test and reinforce comprehension.[11] Our dataset consists of 8 sessions (4 in NO, 4 in O), each consisting of 12 or 16 subjects for a total of 116 participants.[12] Sessions lasted under 90 minutes (including training) and subjects earned on average $14 including a $5 showup fee.

4 Results

In Section 4.1 we discuss some preliminary, general features of the data. In Section 4.2 we report our main, reduced-form results. In Section 4.3 we describe, estimate and interpret a structural model in order to better understand and explain the main patterns in our data.

4.1 Preliminaries

Before reporting our main results, we note some broad characteristics of the data. First, although P1 subjects generally (60% of the time) choose to trade conditionally, they virtually never (less than 5% of the time) choose to trade unconditionally, without regard to signal. The no trade theorem cannot be evalutated for P2 subjects paired with these unconditional traders as it is in

[11] In pilot sessions, subjects expressed confusion over the instructions (particularly in the NO treatment), which appeared to be greatly eased by adding the quiz to the training in our main sessions.

[12] The design uses a nearly identical number of subjects in each treatment: 28 subjects in CURSED and CLEAN and 30 subjects in JOINT and CONFIDENT.

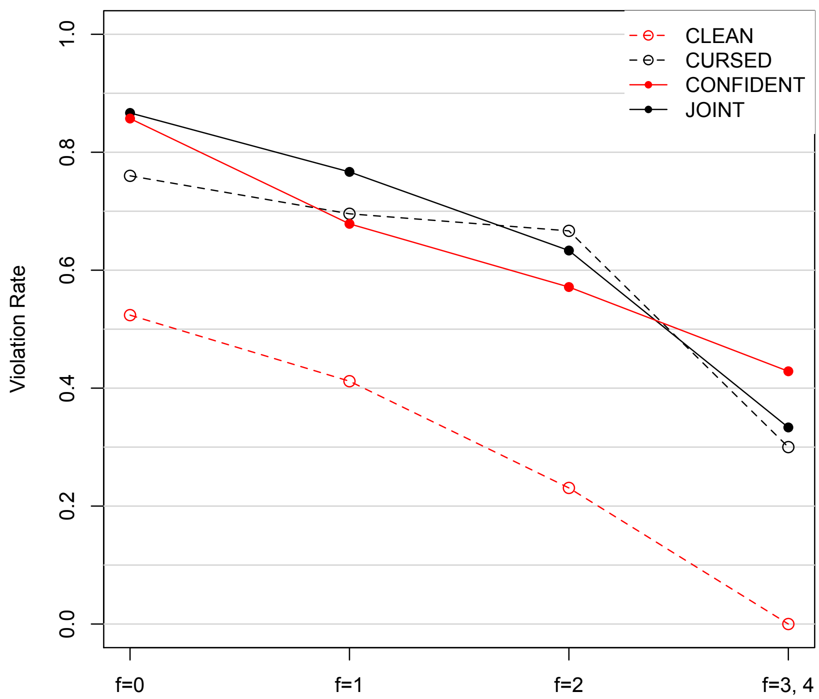

Figure 3: Rates of trade by treatment and fee size.

fact *rational* for P2 subjects to trade in these cases; we drop the small number of P2 subjects in this position in most of the tests below.

Second P1 subjects choose *not* to trade 40% of the time and P2 subjects matched with them are censored from the dataset (particularly, as we will see below, at high fee levels) as they cannot make trading decisions in response. We note, importantly, that this censoring is exogenous to the P2 subjects and therefore doesn't introduce any bias into estimates of P2 behavior.

Third, in the Estimation Stage, subjects correctly estimated the state 65% of the time, with nearly 70% of individual subjects accurately guessing the state 3 or 4 (out of 5) times, suggesting little heterogeneity in accuracy. We find no evidence that subjects were aware of their relative accuracies when making trading decisions – the correlation between accuracy and willingness to trade is in fact negative, but small (-0.048) and statistically insignificant ($p = 0.724$).

4.2 Treatment Level Results

Our main question is how violation of the no trade theorem (as described for our environment in Proposition 1) changes with our treatment variables. To measure this, we examine the rate at which subjects trade when their signal suggests their own asset is less valuable than their potential trading partner's asset. That is, we study the rate at which subjects who are assigned red (blue) items choose to trade under the contingency that their signal (their own estimate, e_i, in O sessions, their guesser's estimate, e_k, in NO sessions) indicates that the state is blue (red). Unless otherwise noted, we will refer to the rate at which subjects are willing to trade in these circumstances simply as the "rate of trade."

We document our treatment level results using two closely related reduced form pieces of evidence. First, in Figure 3 we plot the rate at which subjects choose to trade in each treatment as a function of the fee level. Because we observe few P2 players at higher fee levels (generating as few as 7 observations at a fee of 4 for P2 players), we bin the fee levels at the highest fees (giving us at least 20 data points for each rate shown on the plot). Second, we estimate two statistical models – a linear probability model (LPM) and Logit model – each with random effects at the subject level (note that subjects make all of their decisions without feedback, making subject decisions independent of one another). In each case, the dependent variable is a dummy for trade and the dependent variables include a full set of interactions of our three treatment variables including a dummy for subject type ($P1$, taking a value of 1 if the subject is a P1 type), a dummy for session type (O, taking a value of 1 if the session is an O type session) and fee level, f. Estimates are reported in Table 2.

Our first finding, is that when none of our hypothesized reasons for trade hold – in the CLEAN treatment with high fees – subjects do not violate the no trade theorem at all. This is visible clearly in Figure 3 (the CLEAN series at large fees is equal to zero) and the fit of specification (1) in Table 2 at $f = 4$ (with all other treatment variables set to zero) is indistinguishable from zero by a Wald test ($p = .640$).

Result 1. *When relative overconfidence and cursed reasoning are eliminated and the cost of violating the no-trade theorem is large, subjects do not trade.*

Second, Figure 3 shows that trade is far more likely – 30-40 percentage points higher – at each fee level in every treatment other than CLEAN. The statistical models reported in Table 2 allows us to break down the effects driving this. Both models reveal consistent statistical conclusions, but the LPM model has easier to interpret coefficients, providing estimates of percentage changes

	LPM	Logit
	(1)	(2)
P1	0.290***	1.339**
	(0.106)	(0.634)
O	0.290***	1.369**
	(0.103)	(0.626)
f	−0.145***	−1.158***
	(0.037)	(0.337)
P1 $\times$ O	−0.202	−0.717
	(0.134)	(0.842)
P1 $\times$ f	0.027	0.525
	(0.044)	(0.366)
O $\times$ f	0.010	0.427
	(0.051)	(0.399)
P1 $\times$ O $\times$ f	−0.052	−0.711
	(0.062)	(0.456)
Constant	0.535***	0.408
	(0.082)	(0.477)
Observations	448	448
Log Likelihood	−278.614	−246.138
Akaike Inf. Crit.	577.229	510.276
Bayesian Inf. Crit.	618.277	547.220
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 2: Random effects estimates of the effect of treatment variables on rates of violations of NTT. Specification (1) is a linear probability model (LPM) and specification (2) a logit model.

generated by each of the covariates. Both the $P1$ and O variables are significant, positive and estimated at nearly 0.3, suggesting that cursed reasoning and overconfidence each add about 30 percentage points to the rate of trade. Moreover, the interactions between these variables and the transaction fee, f, are insignificant suggesting that the effect of both of these variables is roughly constant over costs of trade. This provides us with a second result.

Result 2. *Introducing scope for relative overconfidence (the CONFIDENT treatment) or cursed reasoning (the CURSED treatment) each generate about a 30 percentage point increase in rates of trade at each fee level.*

Third, there is a clear effect of the cost of deviating from equilibrium behavior – measured by fee size, f – on rates of trade in all treatments in Figure 3, and the effect is similar across treatments: as the cost of deviating from no-trade theorems rises, rates of violation drop. Estimates in Table 2 confirm this result, with the f variable being estimated significantly negative in each case. Moreover, again, insignificant interactions with the treatment variables $P1$ and O suggest this effect of costs is not impacted significantly by relative overconfidence or cursed reasoning.

Result 3. *As the costs of violating the no-trade theorem rises, the rate at which subjects trade falls and at equal rates across treatments. This results supports the idea of cost-driven noisy deviations from best response as described in the Quantal Response Equilibrium model.*

Finally, Figure 3 shows that rates of trade in the JOINT treatment – where both cursed reasoning and relative overconfidence are possible – is virtually identical to rates of trade in either the CURSED and CONFIDENT treatments alone. Thus there is little evidence of a cumulative joint effect of relative overconfidence and cursed reasoning on trades. Summing O, $P1$ and $O \times P1$ estimates in Table 2 we have an estimated rate of trade in JOINT of 0.378 which is insignificantly different ($p > 0.2$ in all cases via Wald tests) from the estimate of 0.290 shared by the CONFIDENT (measured by the O variable) and CURSED (measured by the P1 variable) treatments, confirming that the effects of relative overconfidence and cursed reasoning do not accumulate. Although we estimate that cursed reasoning and relative overconfidence each generate a rate of trade around 30%, the two combined generate a rate of trade far below a cumulative 60% and in fact not much higher than the 30% estimated in each case individually. We pose this as a next result.

Result 4. *Relative overconfidence and cursed reasoning do not have a cumulative effect on rates of trade. Trade is no higher when both sources of trade are available than when either is available alone.*

4.3 Structural Estimation

Our results show large effects of both relative overconfidence and cursedness on trade. They also suggest there is little cumulative effect when the two are allowed to combine. In this section we structurally estimate key parameters from the model to get a better understanding of what is driving these patterns and to assess the magnitude of cursedness, relative overconfidence and noise in decision making.

4.3.1 Estimation Strategy

We estimate a model that interlaces the quantal response equilibrium model (QRE) with a cursed equilibrium model in a fashion similar to models studied in Carrillo and Palfrey (2009) and Camerer et al. (2016). As in QRE, the probability a subject and his counterpart choose to trade conditional on their respective signals depends on the perceived gains from trading relative to not trading. As in the cursed equilibrium model, each subject has correct beliefs about the distribution of the actions taken by her counterpart but believes that with probability $\chi \in [0,1]$ the action of her counterpart is *independent* of her counterpart's signal. To this we add direct estimation of subject i's' beliefs about the accuracy of her own and others' estimates of the state, allowing us to measure and assess the size and nature of relative overconfidence.

To estimate this model, we allow the probability a subject chooses to trade conditional on a signal to depend on the perceived gain from trading relative to not trading. Using the standard logit specification, the probability player i chooses to trade conditional on signal k is given by

$$\sigma_{i,k} = \frac{e^{\lambda \Delta_{i,k}}}{1 + e^{\lambda \Delta_{i,k}}}. \tag{4}$$

where, as in Section 2, $\Delta_{i,k}$ is the perceived gain from trading, given by:

$$\Delta_{i,k} = \{\tau_{i,k}\left[2(\overline{v} - \underline{v})(\pi_{i,k} - 0.5) - f\right]\} \tag{5}$$

and λ parameterizes the sensitivity of trade to this perceived gain. Under QRE, the conditional probability, $\tau_{i,k}$, a counterpart will trade is

$$\tau_{i,k} = \sum_{l \in \{1,2\}} \sigma_{j,l} Prob(s_j = l | s_i = k)$$

and the probability trade is profitable conditional on the counterpart being willing to trade, $\pi_{i,k}$, is given by

$$\pi_{i,k} = \frac{Prob(s_i = k|\omega_j) \sum_{l\in\{1,2\}} [\sigma_{j,l} Prob(s_i = l|\omega_j)]}{\sum_{h\in\{1,2\}} \left\{ Prob(s_i = k|\omega_h) \sum_{l\in\{1,2\}} [\sigma_{j,l} Prob(s_i = l|\omega_h)] \right\}}$$

To estimate the key cursed reasoning parameter χ, we note that cursed reasoning can only occur for P1 subjects by design (they are the only subjects that must make an inference about their counterpart's trading strategy). For them, the beliefs just discussed, $\tau_{i,k}$ and $\pi_{i,k}$ can be described by the following expressions:

$$\tau_{i,j} = \sigma_{j,i}[\mu_i(1-\mu_j) + (1-\mu_i)\mu_j] + \sigma_{j,j}[\mu_i\mu_j + (1-\mu_i)(1-\mu_j)] \tag{6}$$

$$\pi_{i,j} = \chi\mu_i + (1-\chi)\frac{\mu_i\left[\sigma_{j,i}(1-\mu_j) + \sigma_{j,j}\mu_j\right]}{\mu_i\left[\sigma_{j,i}(1-\mu_j) + \sigma_{j,j}\mu_j\right] + (1-\mu_i)\left[\sigma_{j,i}\mu_j + \sigma_{j,j}(1-\mu_j)\right]} \tag{7}$$

$$\tau_{i,i} = \sigma_{j,i}[\mu_i\mu_j + (1-\mu_i)(1-\mu_j)] + \sigma_{j,j}[\mu_i(1-\mu_j) + (1-\mu_i)\mu_j] \tag{8}$$

$$\pi_{i,i} = \chi(1-\mu_i) + (1-\chi)\frac{(1-\mu_i)\left[\sigma_{j,i}(1-\mu_j) + \sigma_{j,j}\mu_j\right]}{(1-\mu_i)\left[\sigma_{j,i}(1-\mu_j) + \sigma_{j,j}\mu_j\right] + \mu_i\left[\sigma_{j,i}\mu_j + \sigma_{j,j}(1-\mu_j)\right]} \tag{9}$$

P2 subjects, by contrast, cannot suffer cursed reasoning, again by design: they are directly shown that their counterpart will agree to trade if and only if their counterpart receives a signal suggesting trading is advantageous ($\sigma_{j,i} = 1, \sigma_{j,j} = 0$). This affects not only a subject's belief about the likelihood of a trade, but also his posterior belief about the state. In conducting the estimation we therefore restrict P2 subjects to correctly update their beliefs as in a Bayes-Nash equilibrium, generating the following expressions:

$$\tau_{i,j} = \mu_i(1-\mu_j) + (1-\mu_i)\mu_j \tag{10}$$

$$\pi_{i,j} = \frac{\mu_i(1-\mu_j)}{\mu_i(1-\mu_j) + (1-\mu_i)\mu_j} \tag{11}$$

$$\tau_{i,i} = \mu_i\mu_j + (1-\mu_i)(1-\mu_j) \tag{12}$$

$$\pi_{i,i} = \frac{(1-\mu_i)(1-\mu_j)}{\mu_i\mu_j + (1-\mu_i)(1-\mu_j)} \tag{13}$$

In order to estimate the form and degree of relative overconfidence displayed by subjects, we directly estimate the beliefs subjects form about the accuracy of their own (μ_i) and counterparts' (μ_j) beliefs about the state. In doing so, we estimate a model that is flexible enough to allow beliefs about other subjects' accuracy to vary across treatments (which will, as we show below, be important for understanding the results of the experiment). Specifically, we estimate a subject's belief about the accuracy of her own estimate as the parameter $\tilde{\mu}$ and allow her beliefs about the

	(1)	(2)
λ	1.584	1.610
	(0.267)	(0.227)
χ	0.909	0.822
	(0.578)	(0.141)
$\hat{\mu}$	0.518	0.519
	(0.027)	(0.029)
$\tilde{\mu}$	0.628	0.629
	(0.020)	(0.029)
$\overline{\mu}$	0.635	
	(0.080)	
Log Likelihood	-533.165	-533.196

Note: MLE standard errors in parentheses.

Table 3: Maximum Likelihood estimates of the model parameters.

accuracy of others' estimates to vary by treatment: $\hat{\mu}$ is the accuracy the subject attributes to her counterpart in the O treatment (where $\mu_j = \hat{\mu}$) while $\overline{\mu}$ is the accuracy attributed to guessors drawn for the subject and her counterpart in the NO sessions (recall, by design, $\mu_i = \mu_j = \overline{\mu}$ in NO sessions). This feature of our model is important for understanding behavior in the experiment and we can clearly reject the hypothesis that μ_j is constant across treatments (p-value=0.002) using a standard likelihood ratio test (we provide some interpretation of this pattern in our concluding Discussion).

4.3.2 Estimation Results

We estimate the model's parameters, $(\lambda, \chi, \overline{\mu}, \hat{\mu}, \tilde{\mu})$, by maximum likelihood and report results as specification (1) of Table 3. We first consider results on beliefs and relative overconfidence. First, $\tilde{\mu}$ is estimated at 0.628, almost exactly matching subjects' actual estimation accuracy of 0.65 ($p = 0.64$), suggesting that subjects have quite rational assessments of their own beliefs. Second, subjects attribute much lower accuracy to beliefs formed by their own trading partners: $\hat{\mu}$ is estimated at 0.518, meaning subjects trade as though their counterparts' beliefs are essentially white noise (we cannot reject that it is different from 0.5, $p = 0.5$). In the terminology of Eyster et al. (2015), trade in our experiment is thus driven not by *overconfidence* but rather by *dismissiveness*.

Interestingly, subjects only form dismissive beliefs about the accuracy of another subject's guesses when they are actually engaged in speculative (and, perhaps not incidentally, zero-sum) trade with that subject: $\overline{\mu}$ (the accuracy attributed to guessors drawn both for the subject and her counterpart in the NO treatment), is estimated much higher than $\hat{\mu}$ and is almost identical to both $\tilde{\mu}$ and the actual average accuracy of guesses observed in the sample (we consider interpretations of this finding in the concluding Discussion). Because we cannot reject the hypothesis that subjects simply hold correct beliefs of the accuracy of subjects they are not engaged with trading with – that is that these beliefs are parameters of the environment – we estimate a second, restricted model in which $\overline{\mu}$ is simply set to the sample average accuracy (this restriction cannot be rejected using a likelihood ratio test at standard confidence levels). As specification (2) in Table 3 shows, the result of this exercise gives us virtually identical parameter estimates but, importantly, allows us to more precisely estimate the other parameters of the model.

Both structural specifications estimate large values of χ, suggesting that subjects are not only prone to relative overconfidence, but are also prone to quite severe cursed reasoning. In particular, the estimates show that subjects believe their trading partners are over 80% likely to neglect their own signals when deciding whether to trade (when in actuality partners are less than 5% likely to neglect their own signals). This estimate is statistically significant in the more precisely estimated specification, (2).

Do these estimates and their magnitudes cohere with the reduced form results documented in the previous section? Figure 4 (a) plots simulations based on estimates from Table 3 showing rates of trade as a function of treatment and fee level, mirroring Figure 3. The results match those from the data in Figure 3 almost perfectly. In particular they reveal that our estimated parameters are sufficient to generate the large, 30 percentage point increase in trade we observe when both relative overconfidence and cursed reasoning are allowed to influence behavior (and that these effects generally steady over fee levels). Importantly the results also reveal a lack of an additive effect of relative overconfidence and cursed reasoning, with JOINT rates being about the same as both CURSED and CONFIDENT rates of trade, just as in our data.

Result 5. *Structural estimates suggest: (i) subjects are affected by cursed reasoning, believing the action of their counterpart is independent of their counterpart's signal with a probability around 0.8, (ii) subjects in O sessions incorrectly underestimate the accuracy of others' guesses to be around 52%, while they correctly estimate an accuracy for their own guesses of around 63%, and*

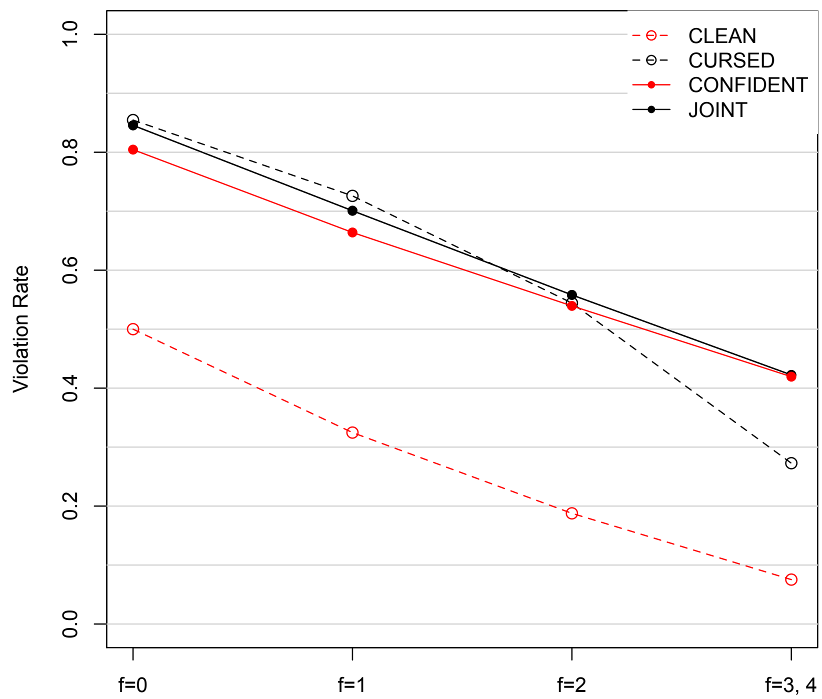

Figure 4: Simulations of rates of trade based on structural estimates.

thus trade due to dismissiveness rather than overconfidence and (iii) these biases are statistically significant and, combined with estimated levels of quantal response noise in decision making, explain the patterns we observe in our experimental data.

An important finding from the structural estimates is that subjects believe their trading partners' estimates of the state are considerably less accurate than they actually are – a form of relative overconfidence that Eyster et al. (2015) refer to as dismissiveness. What's more our estimates suggest that the degree of dismissiveness displayed by subjects is extremely severe: subjects believe their counterparts' efforts to estimate the state are essentially noise and that their beliefs are therefore no better than the prior assigned to state. This severe discounting of the beliefs of trading partners is the reason cursedness and relative overconfidence do not have any cumulative effect in our data. As Eyster et al. (2015) point out (and as we discuss in Section 2.3, above), while overconfidence (belief that your own guess is considerably more accurate than it actually is) can be a *complement* to cursed reasoning, generating more trade than cursed reasoning would generate alone, dismissiveness is a *substitute* for cursed reasoning. Specifically if a subject believes her counterpart's beliefs are no better than random, she will completely ignore the information in her counterpart's signal when choosing whether to trade, generating exactly the same effect of cursed beliefs and via the exact same mechanism. Cursed beliefs cannot therefore inspire additional trade when a subject is dismissive because dismissiveness already has the subject ignoring her counterpart's signal altogether (and vice versa).

5 Discussion

Speculative trade is ubiquitous in naturally occurring markets but understanding its source is difficult using naturally occurring data, where motives for trade abound (e.g. liquidity, risk sharing, insurance, private information). Laboratory experiments can control the set of available motives for trade and recent experiments (Angrisani et al. (2008), Carrillo and Palfrey (2011)) have uncovered evidence that subjects trade (and at high rates) on the basis of private information – a motive ruled out by no-trade theorems (e.g. Milgrom and Stokey (1982), Tirole (1982)). We report results from a diagnostic experiment designed to understand the behavioural causes of these violations of no-trade theorems by systematically removing potential explanations, one-by-one, and examining the resulting impact on trade. This exercise generates stark results: in treatments in which most of these potential explanations are available, subjects are willing to trade more than 70% of the time; in treatments in which none are available, subjects cease to trade entirely.

First, our experiment rules out risk preferences as an explanation in all of our treatments by studying a particularly simple setting in which subjects' only decision is whether to swap equally risky Arrow securities with one another. This trading environment not only rules out risk-based explanations theoretically, it also has the advantage of doing so in a way that makes the lack of insurance motives for trade particularly transparent to subjects. We nonetheless observe subjects trading at very high rates, suggesting that risk preferences are not a primary driver of behavior in these types of experiments (and no driver at all in ours).

Second, we are able to assess the degree to which trade is driven by noisy behavior generated by weak incentives (due, for instance, to bounded rationality or idiosyncratic demand effects favoring trade) by levying a fee, f, on subjects for executing trades. Our game is designed so that the cost of violating no-trade theorems is completely captured by f, making these costs both completely exogenous and highly salient, allowing us to crisply vary the plausibility of this sort of weak incentive explanation. When we raise trading fees to significant levels we remove this sort of noise and are thereby able to eliminate trade altogether in settings where no other channel for trade is available. However, we show that, even at high fee levels, significant trade persists when over-confidence and cursed reasoning (discussed below) remain as motives for trade.

Third, our design allows us to study the degree to which cursed reasoning – systematic under-estimation of potential trading parterns' strategic sophistication (Eyster and Rabin (2005)) – acts as a driver of trade. In some treatments subjects must make their own inferences about their potential trading partners' strategies when deciding whether to trade. In these treatments subjects trade at a significant rate, behaving as though they believe their counterparts systematically fail to condition decisions to trade on private information. In other treatments we directly show subjects how their counterparts plan to condition trades on private information, removing cursed reasoning as an explanation and causing rates of trade to drop substantially as a result. Structural estimates based on these results suggest that subjects believe their trading partners are over 80% likely to neglect private information in making trade decisions when partners are in fact less than 5% likely to neglect their information in this way.

Fourth our design allows us to study whether trade is driven by relative overconfidence about the accuracy of beliefs about the state. When subjects receive their private information by performing a skill-based estimation task (our O treatments), they trade to a significant degree due to relative overconfidence. However when subjects' information instead comes from estimates made by third parties (called "guessers") in our NO treatments, relative overconfidence disappears as an explanation and trade drops considerably as a result. Our cursed reasoning results suggest that subjects under-estimate the sophistication of other traders, and in fact our results on relative over-confidence suggest the same thing: while subjects show little evidence of being overconfident of their own abilities (that is over-estimating the accuracy of their own beliefs), they act as though their counterparts' estimates are little better than white noise (a type of relative overconfidence that Eyster et al. (2015) call dismissiveness). Our design allows us to conclude this because rates of trade are no higher when both cursed reasoning and relative overconfidence are available to influence behavior than when only one or the other is available – a distinctive pattern that is consistent only with an extreme form of dismissiveness in which subjects put virtually no weight on the value of their counterparts' private information.

The specific set of causes for trade identified in our results has implications for the external validity of no-trade violations measured in the laboratory. As Eyster et al. (2015) point out, except in the most extreme cases, the effects of relative overconfidence on trade will tend to diminish as the number of other traders in the market grows large, meaning that relative overconfidence alone will not tend to generate excessive trade in large scale markets.[13] Were (moderate) relative

[13]Though see Daniel and Hirshleifer (2015) for a contrary view.

over-confidence the only cause for trade identified in our data we therefore might expect laboratory no-trade results to fail to scale up in naturally occurring financial markets. However, our results reveal that over-trade is generated both by cursed reasoning and a particularly severe version of dismissiveness that mirrors the effects of cursed reasoning and, as Eyster et al. (2015) argue, these drivers of trade will tend to survive as markets grow large.

Overall, our results show that subjects' willingness to violate no-trade theorems is driven largely by their under-estimation of others' abilities: subjects both underestimate others' ability to form coherent strategies (cursed reasoning) and to make accurate inferences regarding the state (dismissiveness), in each case generating signficant trade. Interestingly, subjects only display dismissiveness towards their potential trading partners: there is no evidence that subjects are dismissive of the beliefs of third party guessers in our NO treatments and in fact our results suggest that subjects form quite realistic beliefs about the accuracy of these guessers. Why do subjects underweight the accuracy of their potential trading partner's beliefs but not those of other subjects? The data does not give us direct evidence but we can speculate on a few possibilities. One is that subjects view counterparts in zero sum interactions as adversaries and naturally form dismissive beliefs about adversaries, undervaluing their abilities (e.g. Klein and Kunda (1992)). This leads to dismissive beliefs for potential trading partners but not for third party guessers. Another is that subjects think better of third party guessors' beliefs than those of their trading partners due to *optimism bias* (e.g. Mraz (2013)): subjects are (weakly) better off when guessors in the NO treatment are accurate but not when their trading partners are accurate, meaning optimism bias might generate high accuracy beliefs in the former case but not in the latter. Regardless, investigating this asymmetry and its potential implications for expectations formation and trade in markets seems an obvious avenue for future research.

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).

Barber, Brad M and Terrance Odean, "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors.," *Journal of Finance*, 2000, *LV* (2), 773–806.

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.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, pp. 861–898.

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.

Caplin, Andrew and Mark Dean, "Revealed Preference, Rational Inattention, and Costly Information Acquisition," Technical Report, NBER Working Paper 2014.

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.

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.

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

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," Technical Report, National Bureau of Economic Research 2015.

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

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

Greiner, B., "Subject pool recruitment procedures: Organizing experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1* (1), 114–125.

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.

Klein, William M and Ziva Kunda, “Motivated person perception: Constructing justifications for desired beliefs,” *Journal of Experimental Social Psychology*, 1992, *28* (2), 145 – 168.

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.

Moore, Don A and Paul J Healy, “The trouble with overconfidence.,” *Psychological review*, 2008, *115* (2), 502.

Mraz, Guy, “The optimism bias: experiment and model,” Technical Report, University of Melbourne 2013.

Nagel, Rosemarie, “Unraveling in guessing games: An experimental study,” *The American Economic Review*, 1995, *85* (5), 1313–1326.

Ngangoue, Kathleen and Georg Weizsacker, “Learning from unrealized versus realized prices,” 2015.

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

_ , “Do Investors Trade Too Much?,” *The American Economic Review*, 1999, *89* (5), 1279–1298.

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.

Stahl, Dale O and Paul W Wilson, “On players? models of other players: Theory and experimental evidence,” *Games and Economic Behavior*, 1995, *10* (1), 218–254.

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

A Appendix A: Proofs

Proof of Proposition 1:

Proof. By contradiction, assume that there is a Bayes-Nash equilibrium where both players choose T for some signal. Specifically assume player i chooses $a_i = T$ if s_i belongs to a non-empty set $S_i \subseteq \{1, 2\}$ and chooses N otherwise. Similarly assume player j chooses $a_j = T$ if s_j belongs to a non-empty set $S_j \subseteq \{1, 2\}$ and chooses N otherwise.

In a Bayes-Nash equilibrium, player i adopts this strategy if and only if $\Delta_{i,k} \geq 0$ for each signal $k \in S_i$. Substituting the expression for $\Delta_{i,k}$:

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

Because we assume utility functions are strictly increasing, it must be:

$$E_{\mu_i,\mu_j}\tau_{i,k}\left\{[\pi_{i,k}u_i(\overline{v})+(1-\pi_{i,k})u_i(\underline{v})]-[\pi_{i,k}u_i(\underline{v})+(1-\pi_{i,k})u_i(\overline{v})]\right\} > 0$$

which can be rewritten as:

$$E_{\mu_i,\mu_j}2\tau_{i,k}\left[u_i(\overline{v})-u_i(\underline{v})\right](\pi_{i,k}-0.5) > 0$$

which in turn simplifies to:

$$E_{\mu_i,\mu_j}\tau_{i,k}(\pi_{i,k}-0.5) > 0 \qquad (14)$$

We have the following expression for $\pi_{i,k}$:

$$\pi_{i,k} \equiv Prob(\omega_j|s_i = k, a_j = T) = Prob(\omega_j|s_i = k, s_j \in S_j)$$

and by Bayes' rule we can write:

$$\pi_{i,k} = \frac{Prob(s_i = k, s_j \in S_j|\omega_j)Prob(\omega_j)}{Prob(s_i = k, s_j \in S_j|\omega_j)Prob(\omega_j) + Prob(s_i = k, s_j \in S_j|\omega_i)Prob(\omega_i)}$$

We have the following expression for $\tau_{i,k}$:

$$\tau_{i,k} \equiv Prob(a_j = T|s_i = k) = Prob(s_j \in S_j|s_i = k) = \frac{Prob(s_i = k, s_j \in S_j)}{Prob(s_i = k)}$$

Using the previous expressions for $\pi_{i,k}$ and $\tau_{i,k}$, noting that $Prob(\omega_j) = Prob(\omega_i) = 0.5$ and $Prob(s_i = k) = 0.5$, condition (14) yields:

$$E_{\mu_i,\mu_j}\left[Prob(s_i = k, s_j \in S_j|\omega_j) - Prob(s_i = k, s_j \in S_j|\omega_j)\right] > 0$$

and since we assume signals are conditionally independent this can be rewritten as:

$$E_{\mu_i,\mu_j}\left[Prob(s_i=k|\omega_j)Prob(s_j\in S_j|\omega_j)-Prob(s_i=k|\omega_i)Prob(s_j\in S_j|\omega_i)\right]>0$$

This can be rewritten as:

$$E_{\mu_i,\mu_j}\left[Prob(s_i=k|\omega_j)\sum_{s_j\in S_j}Prob(s_j|\omega_j)-Prob(s_i=k|\omega_i)\sum_{s_j\in S_j}Prob(s_j|\omega_i)\right]>0$$

Because this condition holds for any $k\in S_i$, it follows that:

$$E_{\mu_i,\mu_j}\sum_{s_i\in S_i}\sum_{s_j\in S_j}\left[Prob(s_i|\omega_j)Prob(s_j|\omega_j)-Prob(s_i|\omega_i)Prob(s_j|\omega_i)\right]>0 \tag{15}$$

By a similar argument, player j agreeing to trade implies that:

$$E_{\mu_i,\mu_j}\sum_{s_i\in S_i}\sum_{s_j\in S_j}\left[Prob(s_i|\omega_i)Prob(s_j|\omega_i)-Prob(s_i|\omega_j)Prob(s_j|\omega_j)\right]>0 \tag{16}$$

But conditions (15) and (16) cannot be satisfied simultaneously. □

Proof of Proposition 2:

Proof. For each player i consider the following trading strategy: choose $a_i=T$ whenever $s_i=j$. Player i adopts this strategy if and only if:

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

that is if and only if:

$$E^i_{\mu_i,\mu_j}\left\{\tau_{i,j}\left[2(\overline{v}-\underline{v})(\pi_{i,j}-0.5)-f\right]\right\}\geq 0 \tag{17}$$

We have the following expressions for $\tau_{i,j}$ and $\pi_{i,j}$:

$$\tau_{i,j}\equiv Prob(a_j=T|s_i=j)=Prob(s_j=i|s_i=j)=\mu_i(1-\mu_j)+(1-\mu_i)\mu_j$$

$$\pi_{i,j}\equiv Prob(\omega_j|s_i=j,a_j=T)=Prob(\omega_j|s_i=j,s_j=i)=$$
$$=\frac{\mu_i(1-\mu_j)}{\mu_i(1-\mu_j)+(1-\mu_i)\mu_j}$$

Substituting the expressions for $\tau_{i,j}$ and $\pi_{i,j}$, equation (17) can be rewritten as:

$$E^i_{\mu_i,\mu_j}\left\{(\overline{v}-\underline{v})\left[\mu_i(1-\mu_j)-(1-\mu_i)\mu_j\right]-f\left[\mu_i(1-\mu_j)+(1-\mu_i)\mu_j\right]\right\}\geq 0$$

Since μ_i and μ_j are independent and with means $E^i\mu_i = \tilde{\mu}$ and $E^i\mu_j = \hat{\mu}$ respectively, this yields:

$$(\overline{v} - \underline{v})\left[\tilde{\mu}(1-\hat{\mu}) - (1-\tilde{\mu})\hat{\mu}\right] - f\left[\tilde{\mu}(1-\hat{\mu}) + (1-\tilde{\mu})\hat{\mu}\right] \geq 0$$

Rearranging this gives:

$$(\overline{v} - \underline{v})\left(\frac{\tilde{\mu}}{1-\tilde{\mu}} - \frac{\hat{\mu}}{1-\hat{\mu}}\right) - f\left(\frac{\tilde{\mu}}{1-\tilde{\mu}} + \frac{\hat{\mu}}{1-\hat{\mu}}\right) \geq 0$$

which can be rewritten as follows:

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

□

Proof of Proposition 3:

Proof. For each player i consider the following trading strategy: choose $a_i = T$ whenever $s_i = j$. Player i adopts this strategy if and only if:

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

that is if and only if:

$$E_{\mu_i,\mu_j}\left\{\tau_{i,j}\left[2(\overline{v} - \underline{v})(\pi_{i,j} - 0.5) - f\right]\right\} \geq 0 \tag{18}$$

The expressions for $\tau_{i,j}$ is the same we used in the proof of Proposition 1, because by definition of cursed equilibrium players have correct beliefs about the distribution of the actions taken by the other players, see Eyster and Rabin (2005). Cursed beliefs affect the posterior probability $\pi_{i,j}$, which now is given by:

$$\pi_{i,j} = \chi\mu_i + (1-\chi)\frac{\mu_i(1-\mu_j)}{\mu_i(1-\mu_j) + (1-\mu_i)\mu_j}$$

Substituting this expression and the expression for $\tau_{i,j}$ into equation (18) we get:

$$\chi 2(\overline{v} - \underline{v})E_{\mu_i,\mu_j}\left\{\tau_{i,j}(\mu_i - 0.5)\right\} +$$
$$+(1-\chi)2(\overline{v} - \underline{v})E_{\mu_i,\mu_j}\left\{\tau_{i,j}\left[\frac{\mu_i(1-\mu_j)}{\mu_i(1-\mu_j) + (1-\mu_i)\mu_j} - 0.5\right]\right\} -$$
$$-fE_{\mu_i,\mu_j}\tau_{i,j} \geq 0$$

The second term on the left-hand side is equal to zero as shown in the proof of Proposition 2. Thus we obtain:

$$\chi 2(\overline{v} - \underline{v})E_{\mu_i,\mu_j}\left\{\tau_{i,j}(\mu_i - 0.5)\right\} - fE_{\mu_i,\mu_j}\tau_{i,j} \geq 0$$

Rearranging gives:

$$\chi \geq \frac{f E_{\mu_i,\mu_j} \tau_{i,-i}}{2(\overline{v} - \underline{v}) E_{\mu_i,\mu_j} \left\{ \tau_{i,j}(\mu_i - 0.5) \right\}}$$

Substituting $\tau_{i,j} = \mu_i(1-\mu_j) + (1-\mu_i)\mu_j$ and taking expectations yields:

$$\chi \geq \frac{f \left(\overline{\mu} - \overline{\mu}^2 \right)}{2(\overline{v} - \underline{v})(\overline{\mu} - 0.5) \left[\overline{\mu} - E\left(\mu^2\right) \right]}$$

□

Proof of Proposition 4:

Proof. The posterior belief of a player who is affected by both cursed thinking and relative over-confidence is:

$$\pi_{i,j} = \chi \tilde{\mu} + (1-\chi) \frac{\tilde{\mu}(1-\overline{\mu})}{\tilde{\mu}(1-\overline{\mu}) + (1-\tilde{\mu})\overline{\mu}} \tag{19}$$

We have the following derivatives:

$$\frac{\partial \pi_{i,j}}{\partial \chi} = \tilde{\mu} - \frac{\tilde{\mu}(1-\overline{\mu})}{\tilde{\mu}(1-\overline{\mu}) + (1-\tilde{\mu})\overline{\mu}} \tag{20}$$

$$\frac{\partial^2 \pi_{i,j}}{\partial \chi \partial \tilde{\mu}} = 1 - \frac{(1-\overline{\mu})\overline{\mu}}{[\tilde{\mu}(1-\overline{\mu}) + (1-\tilde{\mu})\overline{\mu}]^2} \geq 0 \tag{21}$$

$$-\frac{\partial^2 \pi_{i,j}}{\partial \chi \partial \overline{\mu}} = -\frac{(1-\tilde{\mu})\tilde{\mu}}{[\tilde{\mu}(1-\overline{\mu}) + (1-\tilde{\mu})\overline{\mu}]^2} \leq 0 \tag{22}$$

□

B Online Appendix B: Instructions

Instructions were delivered in two parts. We provide instructions for each part for O sessions followed by instructions for each part for NO sessions.

Instructions

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

Your computer screen will display useful information. Remember that the information on your computer screen is PRIVATE. To ensure the best results for yourself, and accurate data for the experimenters, please DO NOT COMMUNICATE with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and the experimenter will come to you.

Economics experiments have a strict policy against deception. If we do anything deceptive, or don't pay you cash as described, then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how the experiment **actually** works and how you earn money in the experiment, and our interest is in seeing how people with an accurate understanding of how their decisions influence their outcomes and earnings make economic decisions.

The Basic Idea.

Items and Colors. This experiment is divided into several **periods**. In each period, you will be given an item that is either blue or red. Every period, the computer will randomly decide (each with equal likelihood) which of the two colors is the **winning color** (items of the winning color have a **value** $14) and which is the **losing color** (items of the losing color have a **value** of $4). Thus either

- The blue item is worth $14 and the red item is worth $4 (if blue wins) OR
- the red item is worth $14 and the blue item is worth $4 (if red wins).

Link Player. Each period another participant in this room will be assigned as your **link player**. If you have a blue item, your link player will have a red item (and vice versa).

Trade. You will decide whether you are willing to trade items with your link player (with the effect that each of you switch the color of your item). You make this decision using an interface like the following one and decide whether to click "Trade My Item" or "Don't Trade My Item" The example on the left is the decision you would make if you were assigned a red item at the beginning of the period and the example on the right if you were assigned a blue item.

○ Trade My Item For Blue	○ Trade My Item For Red
○ Don't Trade My Item	○ Don't Trade My Item

Trade and Transaction Fee. A trade will only occur if both you and your link player decide to trade. If you do, you will both pay a **transaction fee** (to the experimenter) that will vary between 0, 1, 2, 3, and 4 over the periods of the experiment.

Clue of Dots. We will not tell you directly which color is the winner – instead we will give you a **clue** each period about the winner. We will show you (and everyone else) some number of dots on the screen for some number of seconds.

- If the winning color is blue we will show you more blue dots than red dots.

- If the winning color is red we will show you more red dots than blue dots.

Thus, whichever color dots there are more of (blue or red) will be the winning color. When you see the dots, we will ask you to estimate what you think the winning color is and you will earn $4 if you are wrong and $14 if you are right. Just as the winning color will be the same for everyone each period, the number of dots of each color shown as a clue will be **exactly** the same for everyone each period though the arrangement of the dots on the screen will be random and will differ across participants.

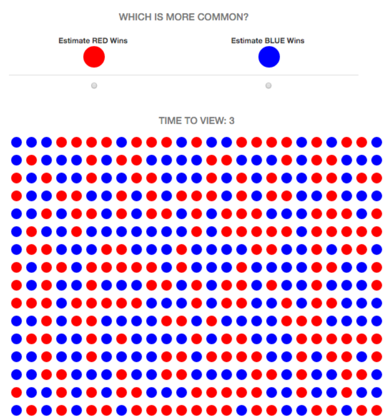

The Trading Table.

Choice Table. Instead of first seeing the clue (the dots) and then deciding whether to trade, the experiment will unfold in the reverse order. You will decide whether you want to trade **contingent** on an estimate of the winning color made later, after everyone sees the clue. To make this contingent decision, you will fill out a **choice table** like the ones below:

There is one column where you provide the decision you would like to make if you (later) estimate that the winner is red and another that you would like to make if you (later) estimate it is blue (it is perfectly fine for the decision to be the same in both cases or different).

2

Determining Trade. **After** both you and your link player have (i) filled out your table, (ii) seen your clue (dots) and (iii) made your guesses about the winning color, the **computer** will implement the trading decisions corresponding to each of your winning color estimates. A trade will happen only if both you and your link player have chosen to trade, given your respective guesses about the winning color.

Earnings.

At the end of the experiment, we will randomly choose one period to be the **pay period** and will pay you based on your decisions that period. We will also randomly determine whether to pay you based on your item's value (minus the transaction fee if you traded) or instead based on the accuracy of your estimate of the winning color in the pay period. We will thus only pay you for based on your item or winning color estimate, but not both.

First and Second Players.

At the beginning of the experiment you will be assigned to be either a **First Player** or **Second Player**:

First Players. The rules given so far are for what we will call **First Players**. If you are a First Player:

- Your link player is another First Player and the two of you are **matched** together (i.e. you are your link player's link player).
- You and your link player make your decisions (fill out your choice tables) at the **same time**, without seeing one another's tables.

Second Players. Some players will instead be **Second Players**. If you are a Second Player:

- Your link player will be a First Player and you will **not** be her link player – instead your link player will have **another** First Player as her link player. Thus, your choices do not affect the earnings or outcomes for any other player, but their decisions affect yours.
- You will make your decision **after** your link player makes her choice and **you will be able to see the decisions from their choice table** before you fill your choice table out.

Choice Tables of Second Players. The images below show what your choice table will look like as a Second Player. Just as with a First Player, the columns allow you to make decisions depending on an estimate of the winning color (made later on). But now, the **table rows** tell you something about what your link player's choice table decisions were.

In the following example, your link player does not want to trade if she estimates that the winning color is red but does if she estimates that it is blue.

MY ITEM IS BLUE

	If I Estimate RED Wins $14 $4	If I Estimate BLUE Wins $14 $4
If My Link Player Estimates RED Wins $14 $4	NO TRADE	NO TRADE
If My Link Player Estimates BLUE Wins $14 $4	Trade My Item For Red Don't Trade My Item	Trade My Item For Red Don't Trade My Item

[Transaction Fee for Completed Trades: **2**]

Your trade decisions will therefore only actually matter in this case if your link player estimates blue is the winner (remember you need both players to want to trade in order for a trade to occur).

On the other hand, the following is an example in which your link player is willing to trade no matter what color she estimates.

Your decisions will thus matter regardless of whether your link player estimates red OR blue. (It is also possible for a link player's decision to be **not** to trade no matter what she estimates.)

One-Way Link. Remember, if you are a Second Player, your decisions will be affected by but will not affect your link player's decisions. And remember, if you are a First Player, your link player will be another First Player (who will not see the results of your choice table before making his or her decision) – your earnings will never be impacted by a Second Player.

Instructions

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

Your computer screen will display useful information. Remember that the information on your computer screen is PRIVATE. To ensure the best results for yourself, and accurate data for the experimenters, please DO NOT COMMUNICATE with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and the experimenter will come to you.

Economics experiments have a strict policy against deception. If we do anything deceptive, or don't pay you cash as described, then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how the experiment **actually** works and how you earn money in the experiment, and our interest is in seeing how people with an accurate understanding of how their decisions influence their outcomes and earnings make economic decisions.

The Basic Idea.

Items and Colors. This experiment is divided into several **periods**. In each period, you will be given an item that is either blue or red. Every period, the computer will randomly decide (each with equal likelihood) which of the two colors is the **winning color** (items of the winning color have a **value** $14) and which is the **losing color** (items of the losing color have a **value** of $4). Thus either

- The blue item is worth $14 and the red item is worth $4 (if blue wins) OR
- the red item is worth $14 and the blue item is worth $4 (if red wins).

Link Player. Each period another participant in this room will be assigned as your **link player**. If you have a blue item, your link player will have a red item (and vice versa).

Trade. You will decide whether you are willing to trade items with your link player (with the effect that each of you switch the color of your item). You make this decision using an interface like the following one and decide whether to click "Trade My Item" or "Don't Trade My Item" The example on the left is the decision you would make if you were assigned a red item at the beginning of the period and the example on the right if you were assigned a blue item.

Trade and Transaction Fee. A trade will only occur if both you and your link player decide to trade. If you do, you will both pay a **transaction fee** (to the experimenter) that will vary between 0, 1, 2, 3, and 4 over the periods of the experiment.

Clue of Dots. We will not tell you directly which color is the winner – instead we will give you a **clue** each period about the winner. We will show you (and everyone else) some number of dots on the screen for some number of seconds.

- If the winning color is blue we will show you more blue dots than red dots.

- If the winning color is red we will show you more red dots than blue dots.

Thus, whichever color dots there are more of (blue or red) will be the winning color. When you see the dots, we will ask you to estimate what you think the winning color is and you will earn $4 if you are wrong and $14 if you are right. Just as the winning color will be the same for everyone each period, the number of dots of each color shown as a clue will be **exactly** the same for everyone each period though the arrangement of the dots on the screen will be random and will differ across participants.

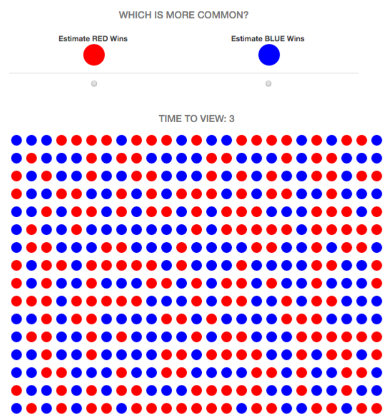

The Trading Table.

Choice Table. Instead of first seeing the clue (the dots) and then deciding whether to trade, the experiment will unfold in the reverse order. You will decide whether you want to trade **contingent** on an estimate of the winning color made later, after everyone sees the clue. To make this contingent decision, you will fill out a **choice table** like the ones below:

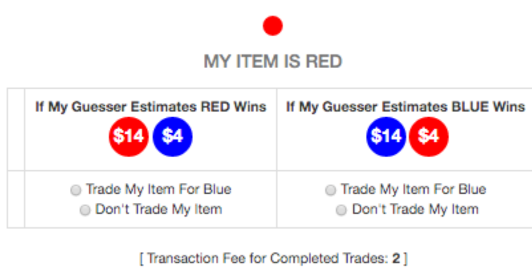

There is one column where you provide the decision you would like to make if the estimate is that the winner is red and another that you would like to make if the estimate is that it is blue (it is perfectly fine for the decision to be the same in both cases or different).

Importantly, **your estimate** will not determine which decision in your choice table is implemented by the computer. Instead,the computer will **randomly select** one other person in this room – referred to as **your Guesser** – and that person's estimate will determine which decision in your choice table is implemented. Thus, when you are filling out your choice table, you will make (i) a decision that the computer will implement if your Guesser estimates red is the winner and (ii) a decision that the computer will implement if your Guesser estimates blue is the winner (the same will be true for your link player).

Determining Trade. **After** both you and your link player have (i) filled out your table, (ii) seen your clue (dots) and (iii) made your guesses about the winning color, the **computer** will randomly select a Guesser for each of you from all of the participants in the room and implement the trading decisions corresponding to each of your Guessors' estimates. A trade will happen only if both you and your link player have chosen to trade, given your respective Guessors' estimates about the winning color.

Earnings.

At the end of the experiment, we will randomly choose one period to be the **pay period** and will pay you based on your decisions that period. We will also randomly determine whether to pay you based on your item's value (minus the transaction fee if you traded) or instead based on the accuracy of your estimate of the winning color in the pay period. We will thus only pay you for based on your item or winning color estimate, but not both.

3

First and Second Players.

At the beginning of the experiment you will be assigned to be either a **First Player** or **Second Player**:

First Players. The rules given so far are for what we will call **First Players**. If you are a First Player:

- Your link player is another First Player and the two of you are **matched** together (i.e. you are your link player's link player).
- You and your link player make your decisions (fill out your choice tables) at the **same time**, without seeing one another's tables.

Second Players. Some players will instead be **Second Players**. If you are a Second Player:

- Your link player will be a First Player and you will **not** be her link player – instead your link player will have **another** First Player as her link player. Thus, your choices do not affect the earnings or outcomes for any other player, but their decisions affect yours.
- You will make your decision **after** your link player makes her choice and **you will be able to see the decisions from their choice table** before you fill your choice table out.

Choice Tables of Second Players. The images below show what your choice table will look like as a Second Player. Just as with a First Player, the columns allow you to make decisions depending on an estimate of the winning color (made later on). But now, the **table rows** tell you something about what your link player's choice table decisions were.

In the following example, your link player does not want to trade if her guesser estimates that the winning color is red but does if her guesser estimates that it is blue.

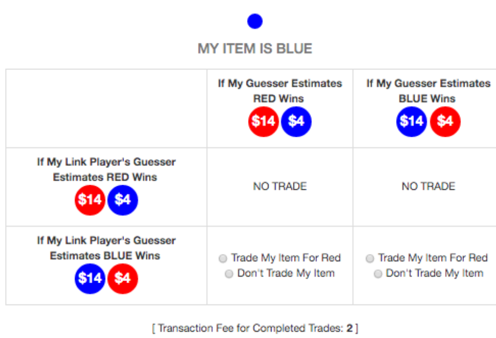
MY ITEM IS BLUE

	If My Guesser Estimates RED Wins $14 $4	If My Guesser Estimates BLUE Wins $14 $4
If My Link Player's Guesser Estimates RED Wins $14 $4	NO TRADE	NO TRADE
If My Link Player's Guesser Estimates BLUE Wins $14 $4	Trade My Item For Red Don't Trade My Item	Trade My Item For Red Don't Trade My Item

[Transaction Fee for Completed Trades: **2**]

Your trade decisions will therefore only actually matter in this case if your link player's guesser estimates blue is the winner (remember you need both players to want to trade in order for a trade to occur).

On the other hand, the following is an example in which your link player is willing to trade no matter what color her guesser estimates.

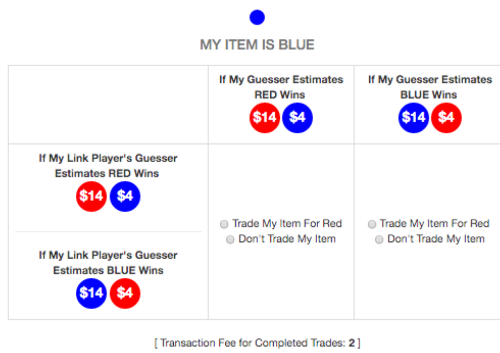

Your decisions will thus matter regardless of whether your link player's guesser estimates red OR blue. (It is also possible for a link player's decision to be **not** to trade no matter what her guesser estimates.)

<u>One-Way Link.</u> Remember, if you are a Second Player, your decisions will be affected by but will not affect your link player's decisions. And remember, if you are a First Player, your link player will be another First Player (who will not see the results of your choice table before making his or her decision) – your earnings will never be impacted by a Second Player.