Trust Me: Communication and Competition in a Psychological Game*

Marina Agranov[†] Utteeyo Dasgupta[‡] Andrew Schotter[§]

February 2, 2023

Abstract

We study, both theoretically and experimentally, a communication game with and without sender competition and embed it in a psychological-game framework where players incur costs for lying, deceiving, and being deceived. We derive the equilibrium predictions of this model and test them in a controlled laboratory experiment. We find that, while the introduction of psychological costs is welfare increasing, the further introduction of competition counteracts this improvement. The latter result is driven by an excessive amount of lying by senders when competition exists and by the inability of receivers to apprehend this dissembling.

1 Introduction

Since the seminal paper by Crawford and Sobel (1982), economists have devoted considerable attention to communication games. These games typically involve an informed sender who sends a message to a less informed receiver, who then takes an action that determines the payoffs to both people. The question investigated in these models is the informativeness of the equilibrium messages sent by senders as a function of the divergence of their preferences from receivers over material outcomes. In these games, senders may lie and deceive others (Kartik (2009), Sobel (2020)) but in the models underlying these games, there is no room for feeling guilt when one misleads others or feeling disappointed when one is lied to. This raises the question of whether the equilibria of these games would be more informative (honest) if senders suffered from both lying and deception aversion (guilt) and receivers could feel disappointment.

This paper aims to answer this question by modeling a market as a sender-receiver game with psychological costs for deception and testing it experimentally in the lab.[1] Implementing such costs can help police these markets, reduce deception, and establish equilibria that are more informative

*We thank the seminar participants at Princeton and Navin Kartik for helpful comments and suggestions.

†Division of the Humanities and Social Sciences, California Institute of Technology. Email: magranov@hss.caltech.edu.

‡Department of Economics, Fordham University. Email: utteeyodasgupta@gmail.com.

§Department of Economics, New York University. Email: andrew.schotter@nyu.edu.

[1] Our theoretical model belongs to the class of psychological games, in which players' payoffs are a function not only of their actions but also of their beliefs.

than the pooling equilibrium that is likely to exist without them. Furthermore, if on top of these deception costs we introduce competition between senders, we might expect a further increase in honesty and efficiency.

What we find both supports and refutes these expectations. While the introduction of psychological payoffs (modeled as a psychological game) does indeed lead to an increase in the honesty of senders and a shift to a more informative equilibrium, the introduction of competition among senders not only counteracts this improvement but almost neutralizes it. Hence the punchline of our paper is simple. While making people psychologically and morally accountable for their actions in markets with psychological costs alters their behavior in a welfare-improving manner, the introduction of competition works in the opposite direction. As we will see this occurs because senders in our laboratory markets react to competition by increasing their lying (for fear that they will lose to their competition if they don't) and receivers fail to account for this increased lying.[2]

We now describe in more detail the main features of our setup, the theoretical considerations about the equilibrium play, and our experimental results. We then discuss the main forces driving these results. We conclude the introduction by summarizing the contribution of our paper and its connection to the literature.

Theoretical Expectations. We study three games. The first game is the standard sender-receiver (seller-buyer) game without competition and without any psychological payoffs. It involves two players: the seller and the buyer. The seller owns one unit of a product that is either of high or low quality and wants to sell it to the buyer. The buyer is interested in purchasing a high-quality product and not the low-quality one. The situation is complicated by the fact that the buyer cannot distinguish between the high- and the low-quality product without actually purchasing it, and instead has to rely on the messages sent by the seller. In this game, players have only material payoffs.

In a second game, in addition to the material payoffs, players receive psychological payoffs, which depend on their types. Although psychological payoffs can be multi-dimensional, we focus on the most prominent ones identified in the literature. Sellers can suffer from lying when they misrepresent the quality of the goods they own, and can also feel guilty if they mislead the buyers about the quality of their goods. Buyers can suffer from disappointment when their expectations about the quality of the goods are misplaced.

Our third game —a game with competition— is identical to our second game except that we add a second seller who competes with the first to sell the good. The competition happens via communication, where each seller sends his own message to the buyer, who then picks one of the sellers based on the received messages.

As we will see, theoretically, the introduction of psychological payoffs to a communication game without competition is unambiguously beneficial. Without such payoffs, the no-trade equilibrium is the only possible equilibrium outcome. However, once psychological forces are introduced, the game without competition admits several informative equilibria in which messages are partially informative, and, as a result, trade occurs with positive probability. All of these equilibria achieve higher expected welfare for both sides of the market than the no-trade pooling equilibrium.

The effect of introducing competition into a game with psychological payoffs is theoretically less clear. Although the game with competition features the same set of equilibria as the game without competition, the theory is silent regarding which equilibrium is more likely to be played. If

[2]This lack of sophistication on the part of receivers is not unique to our paper, however, since it is seen in a number of other papers discussing the disclosure of hidden information (see Jin et al. (2021)).

the same or a more informative equilibrium is played in the presence of competition, buyers stand to benefit from it. However, if sellers with low-quality products feel compelled to lie more due to competition, this will diminish messages' informativeness and negatively affect the buyers' ability to select a better seller to engage in trade. The latter effect can ultimately lead to a selection of a less informative equilibrium in which buyers' welfare is lower. Which result holds is ultimately an empirical question. Consequently, we turn to controlled laboratory experiments to help us sort things out by comparing the results in our three games performed in the lab.

It is worth noting that lying aversion, deriving dis-utility from not telling the truth, is not sufficient to establish all of the equilibria we sustain in our model. This is true because lying, per se, is not a psychological-game force, since the dis-utility from lying does not depend on players' beliefs about others and simply captures the fact that sellers dislike being dishonest, i.e., sending a message that does not match the quality of a product they own.[3] Absent psychological payoffs, which link players' payoffs to their beliefs, both games with and without competition admit at most *one* partially informative equilibrium, which makes our setting considerably less interesting.[4]

Experimental Results. As we mentioned above, while the introduction of psychological costs is welfare improving, their impact is counteracted when competition is introduced since sellers increase their lying which is not responded to effectively by buyers. This leads to higher trade frequency of mostly lower quality goods which hurts both buyers and sellers.

Later in our paper, we explore the mechanisms behind the failure of competitive forces to curb immoral behavior in markets with psychological payoffs. Starting with the sellers, we document that sellers respond to each other by lying more often when they observe another seller lying more often. This dynamic is consistent with the fear of being excluded from interaction with buyers, who tend to choose sellers who claim to have high-quality products. Furthermore, we show theoretically that such a response is optimal when sellers are not sure what strategy their opponent is playing. In other words, the strategic uncertainty about other sellers' actions, which is present both in actual and our laboratory markets, pushes the low-quality sellers towards more lying in order to win the competition for a single buyer.

As for the buyers, we use additional beliefs data collected in the experiment to explore why buyers fail to learn the correct meaning of messages and perceive messages as more informative than they actually are in the markets with competition, and, at the same time, correctly interpret messages in the markets without competition. We contemplate several explanations some of which are cultural and are based on the widespread belief that competition between sellers generally helps consumers, and others are more rational and follow from the simple belief updating when two sellers are present. Indeed, we show theoretically that markets with competition feature a more conservative belief-updating process than the markets without competition; this is driven by the feedback available to buyers in the two markets. Furthermore, we provide empirical evidence of the sluggishness of beliefs in the markets with competition, which is consistent with the theoretical considerations described above.

---

[3]The recent experimental literature has convincingly documented that people possess an intrinsic aversion to lying when messages are cheap talk (Gneezy, 2005; Hurkens and Kartik, 2009; Sanchez-Pages and Vorsatz, 2007) and that people are reluctant to tell even white lies, which benefit both the person telling the lie and the one to whom the lie is told (Erat and Gneezy, 2012).

[4]We discuss in more detail what happens in the model without guilt and disappointment aversion in Section 2.5.

Contribution and Connection to the Literature. Our theoretical model builds on two branches of literature: psychological games that incorporate belief-dependent preferences (Geanakoplos et al., 1989; Battigalli and Dufwenberg, 2007, 2009, 2022) and games with lying costs (Chen et al., 2008; Kartik, 2009; Sobel, 2020). As we mentioned earlier, senders in our model may experience guilt if they lead the receiver on and then double-cross him, and they also can suffer from lying aversion, that is, experience discomfort from lying per se, which does not depend on the receiver's actions or beliefs. The receiver in our model can experience disappointment when she relies on the sender's false claims of a high-quality product and ends up purchasing it. These psychological forces have been identified in the literature as the leading suspects for the communication games with hidden actions and hidden information (Gneezy, 2005; Hurkens and Kartik, 2009; Charness and Dufwenberg, 2006, 2011; Vanberg, 2008; Goeree and Zhang, 2014; Casella et al., 2018; Abeler et al., 2019). Relative to the above-mentioned papers, the contribution of our model is to explicitly model these forces, derive equilibrium predictions for the game they define with and without competition, and document the multiplicity of equilibria that emerges in such a setting.

Our experiment is inspired by the seminal papers of Charness and Dufwenberg (2006, 2011) but uses a design that is conceptually very different. Charness and Dufwenberg (2011) study a different version of a hidden information game and test for the *presence* of psychological forces in communication games by inducing only material payoffs and observing outcomes different from those predicted by the material-payoff-only model. By contrast, we utilize the classical experimental approach pioneered by Smith (1976) to induce psychological types of sellers who differ in their lying and guilt aversion and buyers who differ in their disappointment sensitivity. To be precise, we induce both material and psychological payoffs mimicking our theory considerations and observe how such payoffs affect game outcomes in the presence and absence of competition between sellers.[5] The induced value approach has been successfully implemented in a variety of individual-decision and strategic settings but has not yet been used in games with psychological payoffs. We see our implementation of this approach as one of the contributions of this paper.

The experimental literature concerning the interplay between competition and communication is still in its infancy.[6] The three most closely related papers to ours are Casella et al. (2018), Goeree and Zhang (2014), and Born (2020). Casella et al. (2018) study a communication game with hidden actions and communication among competing senders but do not model the game as a psychological game. The authors find messages are inflated in the game with competition, but these inflated messages induce mostly the same actions from receivers, indicating receivers account for this inflation. Our game is the game with hidden information rather than hidden actions, and our results reveal different patterns: as in Casella et al. (2018), we find a shift in the communication strategies when competition is present, but contrary to Casella et al. (2018), our buyers fail to interpret messages correctly when competition is present. Instead, buyers in our experiment behave as if they believe messages have the same meaning in the presence as well as in the absence of competition.[7]

[5]In Section 3.2, we discuss in detail the challenges associated with implementing the induced value method for psychological payoffs, how we overcome these challenges, and the extent to which the experimenter can successfully control subjects' home-grown psychological costs.

[6]Several studies look at the effects of competition on trust in various environments (Huck et al., 2012; Keck and Karelaia, 2012; Fischbacher et al., 2009). See also Vespa and Wilson (2016), who study experimentally a multi-dimensional communication game with multiple senders and find that in this very different setting, receivers do not use the information optimally.

[7]Similar results are found by Jin et al. (2021) when studying disclosure behavior by sellers in a market. In that market, a failure to disclose the quality of one's product should signal its low quality, yet buyers fail to completely

Closer to our setup, Goeree and Zhang (2014) introduce competition in the hidden-information game studied in Charness and Dufwenberg (2011). They find competition and communication act as substitutes. Communication raises efficiency in the absence of competition but lowers efficiency when competition is present. Similarly, competition raises efficiency without communication but lowers it when parties can communicate with each other. The authors briefly discuss some behavioral explanations that can account for such outcomes, including inequality aversion, guilt aversion, lying aversion, and reciprocity. Although our paper shares some of the features of Goeree and Zhang (2014) with respect to the way we define material payoffs and competition, we take a very different approach by modeling the game as a psychological game in which players exhibit a wide range of emotions (translated into their payoffs). We then obtain theoretical results regarding the effects of competition on market outcomes and players' behavior and test these predictions in a lab experiment in which we induce payoffs associated with these emotions. Despite different approaches, both Goeree and Zhang (2014) and we show competition decreases efficiency in a game with communication.

Finally, Born (2020) studies promise competition between sellers who differ in their intrinsic motivation and costs of breaking promises. This model features both hidden information and hidden actions of sellers. Theoretically, Born shows that, on average, promise competition increases buyers' welfare relative to a no-competition case, because some sellers promise more than they would in the absence of competition. Experimental results reveal that sellers' behavior crucially depends on their game experience as the difference between the competition and the no-competition case was observed only in the first rounds of the experiment. Contrary to Born's results, we observe significant welfare differences between the game with and without competition after subjects have learned to play the game and converged to stable behavior.

Structure of the Paper. We proceed as follows. In Section 2, we introduce our communication game and solve for its equilibria. In Section 3, we describe the experimental design and its implementation. Section 4 contains the results of the experiment, while Section 5 investigates the behavioral mechanisms underlying our results. Section 6 offers some conclusions.

2 The Model

In this section, we present three variants of the communication model, which serve as the basis of our experiment. We first describe our model in the absence of psychological payoffs, that is, using only material payoffs, and then introduce both psychological payoffs and competition.

2.1 The Game without Competition and Material Payoffs only

We study a communication game between an informed seller (he) and an uninformed buyer (she). The seller owns one unit of the product and wants to sell it to the buyer. The product can be either low quality, $q = q_L$, with probability $p > \frac{1}{2}$, or high quality, $q = q_H$, with remaining probability $1 - p$. $Q = \{q_l, q_h\}$ denotes the set of potential product qualities. The seller knows the quality of the product he owns and sends a message, m, to the buyer in an attempt to convince her to purchase his good. Two messages are possible: m_1="The product is really high quality" and m_0 ="The product is low quality." $M = \{m_0, m_1\}$ denotes the set of possible messages. The buyer

adjust for it.

does not know the quality of the good but observes the message sent by the seller. The buyer is interested in purchasing the high-quality product and not the low-quality product. The situation is complicated by the fact that the buyer cannot distinguish the high- from the low-quality good until she purchases it and has to instead rely on the seller's messages. After observing the message, the buyer either buys, or does not, and the game ends.

Figure 1: Material Payoffs in the Game without Competition

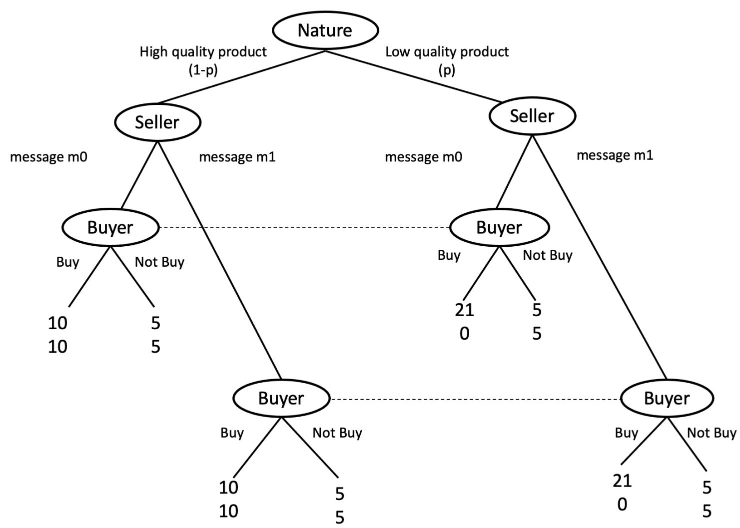

Notes: At each node, the top payoff depicts the seller's payoff, while the bottom one depicts the buyer's payoff. The dashed line indicates the buyer's information set because she does not know the type of seller she is dealing with.

The material payoffs of players are depicted in Figure 1. When a good is not sold, the buyer and the seller each receive a fixed payoff of 5. When a high-quality good is sold, both receive a payoff of 10. The interesting case arises when the seller manages to peddle off a low-quality good: in this case, the seller receives a payoff of 21, while the buyer receives 0. Because in this case, the preferences of the buyer and the seller are misaligned, the potential for lying exists.[8]

Equilibria in the Game without Competition and with Material Payoffs only. Any Bayesian Nash equilibrium outcome in this game features no trade. To see why, assume by contradiction that an equilibrium exists in which, after observing message m_i, the buyer purchases the product with a higher probability than after observing a message m_j. Such behavior is justified if the buyer believes the seller with a high-quality product is more likely to send message m_i than message m_j. However, in that case, the seller with a low-quality product will mimic this behavior and will also send a message m_i, which contradicts our initial presumption. Thus, no equilibrium can exist in which one message entails a higher probability of a high-quality product than another.

[8]We use the term lying to refer to the situation in which the seller's message does not match the product quality he owns, i.e., when a low-quality seller sends the m_1 message and when a high-quality seller sends the m_0 message.

Therefore, the buyer is left with her prior beliefs, and given the material payoffs, *no trade is the only equilibrium outcome.*[9]

2.2 The Game without Competition and Psychological Payoffs

The situation changes when we introduce psychological payoffs. Players are now motivated not only by their material payoffs but also by belief-dependent utilities which are determined by players' strategies and their beliefs. Following the literature, we focus on several psychological forces that have previously been identified as important.[10] We present here the main ingredients of our behavioral model and refer the reader to Section 1 in the Online Appendix for a detailed analysis.

Specifically, we assume the seller may experience guilt and lying aversion, whereas the buyer may experience disappointment. Guilt stems from the fact that a seller can feel bad if he leads the buyer on and then double-crosses her. So, in our game, the seller may feel guilty if he convinces the buyer that he has a high-quality product although he is peddling a low-quality product. The effect of guilt on the seller's payoff depends on players' beliefs. By contrast, a seller may simply experience discomfort from lying whenever he knowingly sends a false message. The disutility from lying does not depend on either how the buyer interprets the message or whether she relied on it for her purchase decision. Finally, the buyer may experience disappointment whenever she relies on the seller's false claims of a high-quality product and ends up purchasing the product.

Formally, we define the seller's psychological type as a pair (L, G), where the first entry is the guilt parameter and the second entry is the lying sensitivity. The psychological type is an innate characteristic of a seller, and as such is known to the seller. The buyer knows the set of all psychological types of sellers, denoted by T^{Seller} and the probability distribution over it, but not the exact type she is dealing with.[11] The seller's decision function maps the product quality he owns into messages for each possible psychological type he may have; that is,

$$s^S : Q \times T^{\text{Seller}} \to M.$$

The buyer's psychological type is captured by a single parameter ω, which denotes the buyer's disappointment sensitivity. The buyer knows her ω but the seller does not and thus has to rely on the distribution of disappointment parameters in the population. We denote the set of all possible psychological types of a buyer as T^{Buyer}. The buyer's decision function maps a message she receives into the purchasing probability for any psychological type she might have; that is,

$$s^B : T^{\text{Buyer}} \times M \to [0, 1],$$

The overall payoffs of players include the material payoffs described above, along with the psychological payoffs determined by players' actions and beliefs. We denote by $b^1_B(m_i)$ the buyer's first-order belief that message m_i is sent by a seller with a high-quality product and by $b^2_S(m_i)$, the seller's second-order belief regarding the buyer's first-order belief about the likelihood that message m_i is sent by a high-quality seller. Then, the overall payoffs of players in this game are

$$\Pi^{\text{Buyer}}\left(s^S, s^B, b^1_B\right) = \Pi^{\text{Buyer}}_{\text{material}} - \omega \cdot \text{Disappointment}$$
$$\Pi^{\text{Seller}}\left(s^S, s^B, b^2_S\right) = \Pi^{\text{Seller}}_{\text{material}} - G \cdot \text{Guilt} - L \cdot \text{Lie}.$$

[9]There are at least two ways to sustain a no-trade equilibrium outcome in our game: one, in which all sellers send message m_0 and another in which all sellers send message m_1.

[10]See Dufwenberg and Gneezy (2000), Gneezy (2005), Battigalli and Dufwenberg (2007), Vanberg (2008), Charness and Dufwenberg (2006, 2011), and Ellingsen et al. (2010).

[11]For simplicity, we omit introducing extra notation here for the probability distribution over sellers' psychological types. The derivation of equilibrium behavior relies on this distribution (see Section 1 in the Online Appendix).

To define the extent of disappointment, guilt, and lying, we use the fact that only two messages are possible and that they have natural meanings and a simple interpretation.[12] Recall that the two messages are $m_0 =$"The product is low quality" and $m_1 =$"The product is really high quality".

We define **lying** as saying things that are not true, that is, misrepresenting your private information. This definition is consistent with theoretical notions used in Kartik (2009) and Sobel (2020) as well as experimental evidence surveyed by Abeler et al. (2019).[13] As Sobel (2020) notes, the definition of a lie depends on the existence of accepted meanings of words. This is exactly what we do in our paper: sellers' messages have precise meanings rather than context-free neutral labels. In our game, lying comes up in two instances: if a seller with a low-quality good sends message m_1, he is lying; and if a seller with a high-quality good sends message m_0, he is also lying, albeit in a way that is typically detrimental to his own causes. The lying parameter L determines the cost one incurs when telling a lie.

By contrast, **guilt** depends both on the message sent by a seller and the players' interpretations of the message (their beliefs). A seller may feel guilty for leading on the buyer (by claiming he has a high-quality product even though he does not) and eventually delivering the low-quality product. This will disappoint the buyer, and the amount of guilt the seller experiences will depend on the buyer's sensitivity to such disappointments. All else being equal, the higher the disappointment parameter ω of the buyer, the more guilty the seller feels when he leads the buyer on and then sells her the low-quality product. Our definition of guilt is reminiscent of the notion of deception in Sobel (2020), where deception depends on how the receiver interprets messages and how her actions might change in response to them.[14] It also relates to Battigalli and Dufwenberg (2007) concept of guilt in games, which captures a failure to live up to receivers' expectations.[15]

Formally, the amount of guilt that the seller experiences is equal to

$$\text{Guilt} = \left(10 \cdot b_S^2(m_1) - 0\right) \cdot \mathbb{E}[\omega|\text{Buyer with type } \omega \text{ buys the product after receiving } m_1],$$

where $10 \cdot b_S^2(m_1)$ represents the seller's belief regarding the payoff that the buyer expects to get when choosing to buy the product after observing m_1, and 0 represents the buyer's actual material payoff when the seller has the low-quality product to deliver after sending m_1. This amount of guilt enters the utility of the seller multiplied by the guilt-sensitivity parameter G.

Finally, the buyer may feel **disappointed** after believing the seller's claim of a high-quality product and buying the low-quality one. We believe the introduction of buyers' disappointment is natural in our game since the fact that buyers may feel disappointed is one way to justify why sellers should feel guilty for leading the buyers on and then double-crossing them.[16] We define

[12] Our theoretical analysis can be extended to incorporate larger message spaces, but the analysis becomes more cumbersome without additional insights.

[13] Abeler et al. (2019) provide a meta-data analysis of experimental work on lying and show that people have a preference to be seen as honest and a preference for honesty per se which explains why people lie less than what theory predicts they should if they only cared about material payoffs.

[14] In the Sobel (2020) framework, the deception and lying capture two very different notions because one player can deceive another without lying, and not all lies constitute a deception.

[15] This theory was coined guilt-from-blame by Charness and Dufwenberg (2011) and used by the authors to explore why communication can be effective in games with hidden information that they study experimentally.

[16] It is, however, not the only way to define guilt in psychological games. In Battigalli and Dufwenberg (2007), theory players experience guilt when they let others down, but this guilt does not depend on how sensitive others are to being disappointed in general. This guilt is determined by the extent to which a player's actions deliver a lower monetary payoff to another player than what the latter expects before the play starts. Our notion of guilt takes an extra step in positing that the extent to which one blames the other for letting her down determines how disappointed she is, and this disappointment is weighted differently by different buyers depending on their sensitivity

this disappointment as equal to the difference between the expected and actual material payoff the buyer receives conditional on observing m_1; that is,

$$\text{Disappointment} = 10 \cdot b_B^1(m_1) - 0.$$

This amount enters the buyer's utility multiplied by the disappointment-sensitivity parameter ω.

Collecting all the terms, we present both psychological and material payoffs in Figure 2. We note that psychological payoffs appear whenever the seller lies and misleads the buyer. For example, take the payoff pair on the bottom right of the game tree where the buyer purchases a low-quality product after being told it was of high quality. Given the message, the buyer thinks the good is of high quality with probability $b_B^1(m_1)$. Although the buyer here expects a payoff of $10 \cdot b_B^1(m_1)$ from purchasing the good, in reality, the good is of low quality, and she ends up with a material payoff of 0. Hence, the magnitude of the buyer's disappointment is $10 \cdot b_B^1(m_1) - 0$, and its effect on her psychological payoffs depends on her realized sensitivity ω, leading to a final payoff of $\omega \cdot (0 - 10 \cdot b_B^1(m_1))$. The seller's payoff is $21 - (10G \cdot b_S^2(m_1) \cdot \omega) - L$. Here, he receives a material payoff of 21 because he managed to peddle off a low-quality good by claiming it to be a high-quality one. However, because this strategy leads to the buyer's disappointment equal to $(10 \cdot b_B^1(m_1) \cdot \omega)$, the seller must subtract a corresponding guilt cost of $G \cdot (10 \cdot b_S^2(m_1) \cdot \omega)$ from the 21, in addition to the lying cost, L, because he lied.[17]

Equilibria in the Game without Competition and with Psychological Payoffs. The introduction of psychological payoffs introduces the possibility of sustaining a **partially informative equilibrium (PIE)** in which some information about the product quality is conveyed in the communication stage. The PIE relies on the existence of psychological utilities that prevent some types of low-quality sellers from mimicking messages of the high-quality sellers.

Specifically, in any PIE, the high-quality sellers send the message m_1 irrespective of their psychological type, while some of the low-quality sellers, those with a relatively low aversion to lying and guilt, might lie. A seller with a low-quality product and psychological type (G, L) prefers to send message m_0 if and only if

$$5 \geq (1 - H[\bar{\omega}(m_1)]) \cdot (5 - L) + H[\bar{\omega}(m_1)] \cdot \left(21 - 10G \cdot b_S^2(m_1) \cdot \mathbb{E}[\omega | \omega \leq \bar{\omega}(m_1)] - L\right).$$

where $H(\cdot)$ denotes the distribution of disappointment-sensitivity parameters of the buyers and $\bar{\omega}(m_i)$ denotes the buyer's type who is indifferent between purchasing and not purchasing the product after observing message m_i. The buyers will choose to buy the product as long as her disappointment sensitivity is not too high, i.e.,

$$\omega \leq \bar{\omega}(m_i) = \frac{2b_B^1(m_i) - 1}{2b_B^1(m_i)(1 - b_B^1(m_i))}$$

The final piece of any PIE is the beliefs' system, which, in equilibrium, must be correct, namely,

$$b_B^1(m_0) = b_S^2(m_0) = 0 \quad \text{and} \quad b_B^1(m_1) = b_S^2(m_1) = \frac{1-p}{1-p+p\cdot\psi},$$

parameter. As a consequence, the seller's guilt is tied to the extent to which he anticipates his actions will disappoint the buyer he is facing.

[17]Because the seller does not know the value of the buyer's ω, the actual payoff to the seller is

$$21 - 10 \cdot G \cdot b_S^2(m_1) \cdot \mathbb{E}[\omega | \text{Buyer with type } \omega \text{ buys the product after receiving } m_1] \cdot 1_{m_i=m_1} - L \cdot 1_{m_i=m_1} \quad \text{if } q = q_L,$$

where $1_{m_i=m_1}$ is an indicator function taking a value of 1 if the seller sends the m_1 message.

Figure 2: Psychological Game without Competition

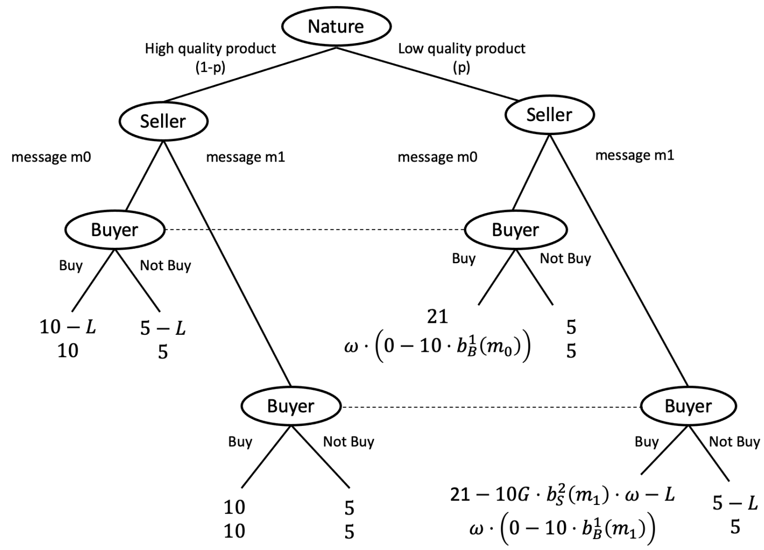

Notes: At each node, the top payoff depicts the seller's payoff, while the bottom one depicts the buyer's payoff. The dashed line indicates the buyer's information set because she does not know the type of seller she is dealing with.

where ψ is the proportion of low-quality sellers who send message m_1 in equilibrium. Note the necessary condition for the existence of PIE is that the proportion of sellers with a low-quality product who lie in equilibrium is not too high. This requirement guarantees at least some buyer types, those with low disappointment sensitivity ($\omega < \bar{\omega}(m_1)$), purchase the product after observing message m_1. The critical task for the buyers after observing an m_1 message amounts to estimating the fraction of sellers who, given the distribution of guilt and lying aversion among them, are reporting truthfully.

2.3 The Game with Competition and Psychological Payoffs

This game with competition has three players: the buyer, seller 1, and seller 2. At the outset of the game, nature draws psychological types of players and product quality for each seller: $\omega \in T^{\text{Buyer}}$ for the Buyer; $(G_i, L_i) \in T^{\text{Seller}}$ for seller i, where $i \in \{1, 2\}$; and $q_i \in Q$ for $i \in \{1, 2\}$. Players' types are private information and are drawn independently from the same distributions as in the game without competition. After observing their psychological types and the quality of their goods (each of which has a probability p of being a low-quality product), both sellers simultaneously submit their messages to the buyer: $(m_i^{\text{Seller 1}}, m_i^{\text{Seller 2}}) \in M \times M$, where $M = \{m_0, m_1\}$. The buyer observes both messages and chooses one of the sellers with whom she will proceed to play the tree game described in Figure 2. The buyer then chooses either to purchase the product or not, and the chosen seller and the buyer receive payoffs specified in the tree game given their strategies. The seller who was not selected by the buyer receives zero payoff.

Equilibria in the Game with Competition and Psychological Payoffs. We focus on symmetric equilibria, in which both sellers use the same communication strategy when they are of the same psychological type, own a product of the same quality, and hold the same beliefs. In general, the game with competition admits the same types of equilibria as the game without competition. These are partially informative equilibria in which the messages are somewhat informative in the sense that sellers who own low-quality products and have high sensitivity to lying and guilt prefer to be truthful and send message m_0. In any PIE, if the buyer receives two different messages from the sellers, she selects the one who sent message m_1, and depending on her disappointment sensitivity ω and her belief $b_B^1(m_1)$, she either buys the product or not. If the buyer receives two m_0 messages, she randomly selects one of the sellers and does not purchase the product. Finally, if the buyer receives two m_1 messages, she selects randomly one of the sellers and purchases the product from him with some positive probability depending on her disappointment parameter.

The exact set of equilibria in games with and without competition depends on the parameterization of the game. In the next section, we discuss our chosen parameters, which are the same parameters we use in our experiment. We chose these parameters so that the two games both have the same sets of equilibria: two PIEs, which differ in the fraction of low-quality sellers who lie in the communication stage.

2.4 Parameterization

To bring our model to the lab, we need to set parameters to be used in all games. We chose the following values. The probability that the product quality is low is $p = 60\%$, which ensures that, absent psychological utilities, the unique equilibrium outcome is the one in which the buyer never purchases the product. The buyer's disappointment parameter ω comes from a uniform distribution $H(\omega) = U[0, 1]$. Conditional on the product quality, four psychological types of sellers are equally likely: $(L, G) \in \{(0, 0), (0, 6), (20, 0), (20, 6)\}$.[18] This reflects the fact that in our experiment, some sellers are motivated by both guilt and lie aversion, some experience only one of these two costs, and some do not suffer from any psychological costs at all. These psychological types are uncorrelated with the quality of the product that a seller owns, and hence represent the seller's internal sensitivity to lying and guilt.

For the experiment and for the analysis presented below, we abstract away from penalizing the sellers who lie and send the message m_0 when they have a high-quality product. Although sending an m_0 message when $q = q_H$ is indeed a lie, it is a self-destructive one and a weakly dominated action. In this case, we do not deduct lying costs for sellers and assume that if $q = q_H$, then $\Pi^{\text{Seller}}(m_0, \text{Buy}) = 10$ and $\Pi^{\text{Seller}}(m_0, \text{Not Buy}) = 5$.[19]

As noted earlier, we chose our parameters in such a way that the set of equilibria that can be supported with and without competition are identical in the games. Specifically, in both games, three equilibria exist:

1. **Pooling equilibrium.** In this equilibrium, all types of sellers send message m_0 and the buyer treats messages as uninformative and does not update her prior beliefs about the product

[18] The discreteness of the psychological-types space of the seller is not a crucial assumption. We use it for simplicity and because it facilitates the comparison between games with and without competition.

[19] The consequences of this assumption is that both games with psychological payoffs admit an additional pooling equilibrium, in which no trade occurs. To sustain this outcome, sellers with both low- and high-quality products must always send message m_0 in the communication stage, and a deviation to the m_1 message should not be interpreted by the buyer as a signal of high quality. Note that this pooling equilibrium is quite fragile as it requires all sellers to send message m_0 to utilize the fact that sending message m_0 is a free lie for a high-quality type.

quality, regardless of the observed message. That is, after observing either message, the buyer believes the chance that she is facing a low-quality seller is 60%. In the game with competition, the buyer randomly selects one seller. In both games, the buyer does not purchase the product and collects a payoff of 5.

2. **PIE1.** In this equilibrium, sellers with a low-quality product and with psychological types $(0, 0)$ and $(0, 6)$ lie and send message m_1 in equilibrium, whereas the remaining sellers with low-quality products truthfully reveal the quality of their products. In the game with competition, the buyer selects a seller with message m_1 if she receives two different messages; otherwise, she randomly selects one seller. In this equilibrium, if the message of the chosen seller is m_1, the buyer believes that there is a 57% chance that this message comes from the high-quality seller and only buyers with relatively low disappointment sensitivity buy the product. Specifically, buyers purchase the good with a probability of 0.51 after receiving an m_1 message from the (chosen) seller. If, however, the chosen seller's message is m_0, the buyer knows this message is sent by the low-quality seller and does not buy the product. The buyer's expected payoff is 5.22 in the game without competition, and 5.29 in the game with competition.

3. **PIE2.** In this equilibrium, only the low-quality sellers with the psychological type $(0, 0)$, that is, those who do not suffer from either lying or guilt, lie and send m_1 in equilibrium. The remaining types truthfully reveal their product quality. In the game with competition, the buyer selects a seller with message m_1 if she receives two different messages; otherwise, she randomly selects one seller. If the chosen seller's message is m_1, the buyer believes the chance that this message comes from the high-quality seller is 73% and this belief is high enough that even the buyer with the highest level of disappointment, $\omega = 1$,, prefers to buy the product. Therefore, after observing message m_1, all buyer types purchase the product. However, if the chosen seller's message is m_0, the buyer knows for sure that the good is of low quality and thus does not buy the product. The buyer's expected payoff is 6.04 in the game without competition, and 6.46 in the game with competition.

The three equilibria described above are ranked in terms of how much information sellers transmit in the communication stage. Define the informativeness of an equilibrium as the difference between buyers' posterior beliefs after observing the two messages, i.e., $Eq^{\text{info}} = b_B^1(m_1) - b_B^1(m_0)$. The larger this difference, the more information the buyer learns from the sellers' messages. Then, the least informative equilibrium is the pooling one, while the most informative one is the PIE2, in which only the low-quality sellers with no psychological costs lie in the equilibrium

$$0 = Eq^{\text{info}}_{\text{POOL}} < Eq^{\text{info}}_{\text{PIE1}} = 0.57 < Eq^{\text{info}}_{\text{PIE2}} = 0.73$$

Informativeness of equilibria directly translates into buyers' expected payoffs: the more information the buyer receives from the seller's messages, the better purchasing decisions she can make. This can be seen by comparing the buyers' expected payoffs across three equilibria holding fixed the presence or absence of competition between sellers. Finally, we note that if the same PIE is played in both games with and without competition, then buyers benefit from competition and earn higher expected payoffs. This happens because the presence of two sellers increases the likelihood that the buyer will select a high-quality seller to play the tree game. However, this observation relies on the assumption that the same equilibrium is played in the two games. Whether such assumption is reasonable or not is ultimately an empirical question, which we address next in our experiment.[20]

[20] Schotter et al. (1996) have shown competition can have an impact on rejection behavior in ultimatum games

2.5 Discussion of Modeling Choices

Before we turn to the experiment, let us discuss some of our modeling choices and their relation to the existing literature.

Game tree and material payoffs. The structure of our game is closely related to the communication game of Crawford and Sobel (1982) in which an informed sender (the seller) sends a message to an uninformed receiver (the buyer), who then takes an action that affects both players' payoffs. We chose to focus on a simpler version of such a game with senders (sellers) possessing high- or low-quality products and receivers (buyers) having to decide whether to buy.

The important feature of our setting is that a buyer knows neither the quality of the seller's product nor the psychological type of the seller she is dealing with; that is, our game belongs to the class of communication games with hidden information. One of the key experimental papers in this literature is that of Charness and Dufwenberg (2011), CD-11 hereafter. The same CD-11 game is used in a follow-up paper by Goeree and Zhang (2014), who introduce competition between sellers and find that communication and competition act as substitutes. Although our game shares some similarities with the game in the above-mentioned papers, important differences also exist.[21] First, in the CD-11 game, the seller can choose not to trade with the buyer even if the buyer wants to trade; in this case, both get a fixed no-trade payoff. This adds an element of reciprocity on the part of the informed player (the seller). By contrast, our game is a pure communication game in which a seller can only send a message to a buyer, and otherwise has no action to take. Second, in the CD-11 game, only high-quality-product sellers prefer to trade, whereas the low-quality-product sellers prefer a fixed no-trade payoff. Our game differs, in that all sellers want to trade irrespective of the quality of their product. Third, the CD-11 game has an additional element that is absent in our game, namely, a positive probability that trade may be prevented from occurring even if both parties agree to trade. This feature adds another layer of uncertainty and reduces the buyer's ability to infer the seller's product quality even at the end of the game. We chose to study what we feel is the simplest communication game with hidden information, which is amenable to psychological payoffs and the introduction of competition among sellers.

The role of psychological payoffs. One may wonder whether a psychological game is at all necessary to study the effect of competition in our communication game. In other words, what would happen in a standard cheap-talk game, in which sellers experience only aversion to lying but not the guilt of misleading the buyers and buyers do not feel disappointed when they expect outcomes that do not materialize? Such a game would be the standard game since sellers' guilt and buyers' disappointment are what make our game a psychological game, i.e., the game in which payoffs depend on players' beliefs. Contrary to that, lying aversion does not depend on sellers' beliefs and captures the pure moral dis-utility from being dishonest.[22]

in which receivers are more willing to accept low offers if the person making such an offer had to compete in a tournament-like setting.

[21] Here, we focus on the one game studied in CD-11, which is closest to our game in the sense that the low-quality seller can gain materially from trading with the buyer if and only if he fools the buyer into buying his product, which the buyer would prefer not to buy. The authors also study another version of the communication game in which the low-quality seller can benefit materially from trading with the buyer, even if the buyer knows she is buying a low-quality product. The introduction of communication between the seller and the buyer is more effective in the second game than in the first one.

[22] We refer the reader to Footnote 3 in the introduction which summarizes the empirical evidence of lying aversion.

This raises the question of what would happen in a restricted version of our model where only lying aversion (and not guilt aversion) exists.[23] Depending on the distribution of the sellers' lying aversion, one can still support a partially informative equilibrium in such a game, in which all high-quality sellers and some low-quality sellers, those with a relatively low aversion to lying, send message m_1, while those with a low-quality product and a high aversion to lying tell the truth. If the fraction of low-quality sellers who lie is not too large, the buyer would choose to buy the product that is advertised as high quality. However, this would be *the only PIE* that can be sustained in such a setting.

The presence of both guilt aversion and lying aversion is what allows us to sustain *two different PIEs* in both games with and without competition and opens up the possibility of competition being welfare decreasing by creating the possibility that it leads to the selection of the welfare-inferior equilibrium. To see this point, re-examine the condition that guarantees that a seller with a low-quality product and a psychological type (G, L) is willing to reveal his type and send m_0 as opposed to m_1 in our game which features all three forces, i.e., the guilt, the lying aversion, and the disappointment aversion:

$$L + 10G \cdot b_S^2(m_1) \cdot \mathbb{E}\left[\omega|\omega \leq \bar{\omega}(m_1)\right] \cdot H[\bar{\omega}(m_1)] \geq 16 \cdot H[\bar{\omega}(m_1)]$$

where $\bar{\omega}(m_1)$ is the buyer type that is indifferent between buying the product and not after observing message m_1. As apparent from this condition, both guilt and lying aversion prevent some low-quality sellers from lying to a buyer. However, the two forces play out differently. The first term on the left captures that low-quality sellers with a high aversion to lying per se are likely to be truthful in the communication stage because their own intrinsic penalty for lying is too high. The second term on the left captures the dis-utility from lying resulting from the guilt that one experiences about misleading the buyer; the extent of sellers' guilt depends on the sellers' beliefs about buyers' interpretation of message m_1. Either of the two channels can prevent low-quality sellers from sending an untruthful message m_1. But, the presence of *both* channels is what allows us to sustain two different PIEs, which can be ranked in terms of how much information is contained in message m_1. Absent the psychological forces, our game would be considerably less interesting because, as mentioned above, there would be no room for competition to select an equilibrium different from the one existing with only lying aversion.

3 Experimental Design

The experiment was conducted in the experimental lab of the Center for Experimental Social Science (CESS) at New York University. We recruited 179 subjects via E-mail from the general undergraduate population at NYU for an experiment that lasted approximately one hour and 45 minutes. Subjects received a show-up fee of $7 and on average received a final payment of $29.50 for their participation. The program used in the experiment was written in Z-Tree (Fischbacher (2007)). We present our experimental design and treatments' variation in Section 3.1. In Section 3.2 we discuss the benefits and challenges of inducing psychological costs in the lab experiment and describe how we deal with eliciting both subjects' actions and subjects' beliefs.

[23] Kartik (2009) analyzes the general version of such a model without competition between sellers and finds that partially informative equilibria exist, in which sellers with higher types pool together, while lower types separate. These equilibria are characterized by inflated language in the sense that sellers might claim they have a higher type than they actually have.

3.1 The Design

Our experiment is a direct implementation of the model described above. We conducted three separate treatments: a Monetary treatment, a No Competition treatment, and a Competition treatment. One unique innovation of our experiment is that we induce the psychological payoffs described in Figure 2 for the two last treatments. So we impose costs on the seller whenever he lies to the buyer and disappoints her. We also impose the disappointment costs on the buyer when she is misled by the seller. As discussed above, these lying, guilt, and disappointment costs (L, G, and ω) are induced and take on different values depending on the player type, in contrast to other experiments in which such costs are typically inferred. Note here that inducing psychological payoffs is no different from inducing material payoffs or risk attitudes, a common practice in laboratory experiments and one of its strengths. If our subjects attempt to maximize their payoff in the experiment, they would be acting as if they had psychological payoffs. Thus, inducing and controlling such psychological payoffs is a fair way to test predictions of psychological games, which is what we do in this paper.[24]

Each experimental session consisted of only one of the three treatments. Once in the lab, subjects were randomly assigned to play the role of either a buyer or a seller, and these roles remained fixed during the entire session. We refer the reader to Section 2 in the Online Appendix for the complete set of instructions in one of the treatments and describe below the main features of the experimental protocol.

In the **No Competition treatment**, the subjects play the communication game described in Figure 2 with the parameters described in Section 2.4. Specifically, the seller's task is to specify a decision function that maps his psychological type and the type of good he is endowed with into a message from the set of messages $\{m_{0,}m_1\}$. The sellers enter their decisions by filling out a table presented in Figure 3.

The buyer's task is to enter a purchasing decision conditional on the message she receives and her sensitivity type ω. Buyers do that in the experiment by entering two cutoff values, $\omega'(m_0)$ for message m_0, and $\omega'(m_1)$ for message m_1, such that whenever the realized value of ω is less than $\omega'(m_0)$ ($\omega'(m_1)$), the buyer buys the good. For values of ω above the cutoff value, the buyer does not buy the good. This decision function essentially suggests buying the good as long as the buyer is not too sensitive to the potential disappointment that stems from being lied to.

In addition to specifying their strategies, subjects are also asked to enter their beliefs. Each buyer is asked to enter a number between 0 and 100 representing her belief that the sellers who sent the message m_i possessed a high-quality good. They did so for both messages m_0 and m_1. Each seller was asked to enter a number representing his (second-order) belief about what the seller thought the first-order belief of the buyer was, upon receiving either message m_0 or m_1.

Once the subjects had specified their strategies and beliefs, these choices were simulated for 10 periods, where the computer randomly determined the quality of the good and a type for each seller, and a sensitivity parameter for each buyer for every period. Further, using the strategies they entered, the computer determined payoffs for them for each of the 10 periods. We call these 10 periods *a block*, and each treatment had 10 such blocks. After each block, the subjects were given time to review their actions and payoffs for the preceding 10 periods before entering their strategies and beliefs again for the next block that determined their payoffs for the next 10 periods. In each

[24]Given that we induce psychological payoffs, our focus is not on assessing whether real-world agents suffer from guilt, disappointment or lying aversion, but rather how their behavior changes in the presence of such psychological motives. By inducing them, we can observe whether behavior in the face of these motives is consistent with what our model predicts.

Figure 3: The interface of Seller's Task

block, subjects maintained their roles but were randomly assigned a new partner.[25]

We use this block design because entering a strategy, a set of beliefs, and reviewing feedback is a time-consuming process, and, hence, would be practically impossible for subjects to do for, say, 50 periods. Our design allows subjects to maximize the amount of feedback they get while economizing on the time they spend mechanically entering their strategies and beliefs. More importantly, we feel this approach is the correct way to conduct experiments using the strategy method, because once a strategy is entered, one might as well receive a lot of feedback on it before being asked to change it.[26] Entering a strategy and receiving only one period of feedback does not allow a subject to learn very much about it.

To determine a subject's payoff in the experiment, we randomly chose one of the 10 blocks, and in that block paid subjects either for their payoffs in the game or for their elicited beliefs, using a quadratic scoring rule (for a similar approach, see Nyarko and Schotter (2002)). Eliciting both actions and true beliefs in psychological games is tricky because payoffs are a function of beliefs. We discuss this issue in detail in Section 3.2 and describe how we dealt with it.

Finally, at the end of the session, we administered two risk-elicitation tasks using the Gneezy and Potters (1997) methodology. In each of these two tasks, we asked subjects to allocate 200 points (translating into $2) between a safe investment, which had a unit return (i.e., returning point for point), and a risky investment, which with probability p returned R points for each point invested and with probability $1 - p$ produced no returns for the investment. In the first task, $p = 0.5$ and

[25]The screenshots depicting feedback that subjects received at the end of each block are presented in Section 3 of the Online Appendix.

[26]Dal Bo and Frechette (2019) use a similar method when they study infinitely repeated prisoners' dilemma games and use the strategy method.

$R = 2.5$, whereas in the second task, $p = 0.4$ and $R = 3$. One of these two risk tasks was randomly chosen to account for payment and earnings from the risk-elicitation task was also added to the earnings from the main task. Conducting two similar tasks with different parameters allows us to reduce measurement errors as shown in Gillen et al. (2019).

In the **Competition treatment**, all procedures were identical to the No Competition treatment, except we had two sellers competing for a single buyer. Hence, the buyer needed to indicate which seller she would buy from given the messages received from each. Four different scenarios could occur: either both sellers sent message m_0, or both sellers sent message m1, or seller 1 sent m_0 and seller 2 sent m_1, or seller 1 sent m_1 and seller 2 sent m_0. For each of these four cases, the buyer specified the probability, a number between 0 and 1, that she wants to be matched with seller 1 (with the remaining probability she was matched with seller 2). Sellers who were not matched were paid zero, whereas those who were matched received payoffs identical to those specified in Figure 2 conditional on their specified strategy and that of the buyer. We again used the block structure for payoffs here and paid either the game payoffs or the belief payoffs for one randomly selected block. In each treatment, payoffs were calibrated so that the payoffs received from beliefs were comparable to those from the game.

In the **Monetary treatment**, although all the procedures were identical to the No Competition treatment, the payoffs did not reflect the psychological costs. Instead, the participants simply played the game with payoffs described in Figure 1. So, a seller was asked to specify the message that would be sent to the buyer for each possible product quality he might possess, and a buyer was asked to specify her purchasing decision for each of the two messages she could receive from the seller. We also elicited buyers' and sellers' beliefs as before. Our experimental design is summarized in Table 1.

Table 1: Experimental Design

Treatment	Number of sessions	Number of subjects
Monetary	3 sessions	58 subjects: 29 Buyers and 29 Sellers
No Competition	3 sessions	52 subjects: 26 Buyers and 26 Sellers
Competition	4 sessions	69 subjects: 23 Buyers and 46 Sellers

3.2 Discussion of experimental design choices

Inducing guilt and lying aversion in the lab. Psychological games take their name from the fact that decision-makers may be affected by their beliefs about others and their beliefs about others' beliefs about them (second-order beliefs). These beliefs can create a variety of emotions on the part of the decision-maker, which would affect how one plays the game. Hence, to properly test a psychological game in the lab, these emotions must be controlled or inferred ex-post given subject behavior.

In this paper, we take the first route and induce guilt and lying aversion by penalizing subjects for lying and misleading others. We do so in line with the standard notion of induced value as originated by Smith (1976), in which an experimenter assigns payoffs to outcomes in such a way that any subject whose utility function is monotonic in lab payoffs will act as if they are maximizing the induced utility function. However, inducing guilt, disappointment, and lying aversion is tricky because people walk into the lab with their own homegrown attitudes toward lying and deceit, and these attitudes may be overlaid on top of or exceed the penalties we impose. This could imply a lack of control.

Our experimental design addresses these concerns and allows us to test whether inducing psychological costs works. To do this note that from a theoretical point of view, either the penalty we impose for lying, disappointment, and guilt is binding or it is not. What we mean by binding is that either the imposed penalty is more severe than the one subjects would impose on themselves given their homegrown attitudes or it is less severe. If it is more severe, we are in control of the subjects' behavior, because our penalties are sufficiently large to be the determining factor in subjects' calculations. Using the language of Smith (1976), this means the Dominance Principle is satisfied; that is, the reward medium dominantly determines changes in the subject's utility. The other case to consider is when a subject's moral aversion to guilt, disappointment, and lying is greater than the penalties we impose. Our experimental design allows us to detect this case by comparing behavior in the Monetary treatment with that in the No Competition treatment, which differs by the inclusion of psychological payoffs in the latter case. Specifically, if subjects had greater resistance to lying or misleading others than the one we imposed in the No Competition treatment, we should observe the same amount or strictly less lying in the Monetary treatment than in the No Competition treatment. The difference between lying in these two treatments would indicate the degree to which induced costs crowd out the homegrown psychological costs.

In the Results section, we address this point by comparing sellers' behavior in the Monetary and No Competition treatment. We show that without induced psychological costs, sellers lie to a far greater extent than they do when such costs are induced (Table 3), providing validation that our technique increases experimental control over psychological forces central to our behavioral model.

Finally, the comparison between the No Competition and the Competition treatments remains valid regardless of whether self-imposed costs are larger or smaller than those induced in our experiments. The reason is that we use the same experimental technique in both treatments and have no reason to believe self-imposed psychological costs should respond to the number of sellers.

Eliciting beliefs and actions in psychological games. In psychological games, payoffs are a function of both actions and beliefs. Therefore, to properly test psychological games in the lab, one needs to elicit subjects' actions and their beliefs about the actions of others. However, asking subjects to report beliefs has two effects: first, it affects subjects' payoffs in the belief elicitation exercise itself, and, second, it affects their payoff in the game, since game payoffs depend on stated beliefs. The second effect implies that subjects might have a strong incentive to report zero beliefs and claim to have no guilt or no disappointment. This will increase a subject's payoff in the tree game since such feelings are subtracted from their material payoffs.[27] In other words, if one takes psychological games seriously and wants to test their equilibrium in the lab, this is a severe and inescapable problem.

In our experiments, subjects are paid both for the beliefs they state and for the actions they choose. In particular, subjects are paid either for their performance in the game or for the accuracy of their beliefs in one randomly selected block of the experiment.[28]The beliefs are elicited using the

[27]In our setup, a seller can inflate his game payoff by reporting a zero second-order belief about message m_1 to minimize the guilt payoff. A buyer can do the same by reporting a zero first-order belief about m_1 to minimize the disappointment payoff.

[28]This randomization diminishes hedging motives which lead to stating beliefs that are not compatible with one's actions (see Costa-Gomez and Weizsacker (2008)).

quadratic scoring rule which penalizes them based on the difference between their stated belief for m_i and the actual observed probability that message m_i comes from a high-quality seller. Similarly, the sellers are penalized based on the difference between the belief they state and the belief buyers state for the same message m_i. We set payoffs for this quadratic scoring rule so that there is practically no incentive to report false beliefs by making the subjects indifferent at the margin between reporting false beliefs or not.[29] Theoretically, this means that they trade off their payoff in the game versus their payoff in the belief elicitation procedure. To help subjects comprehend this, we explained this to them and also told them they had no incentive to report beliefs falsely if they want to maximize the expected dollar payoff in the experiment. This method of announcing to the subjects facing a complicated belief elicitation procedure that truth-telling is the optimal thing to do is the most effective way to elicit true beliefs in the lab as shown in the recent influential paper by Danz et al. (2021) and is commonly done in the experiments.

We now turn to our data to investigate whether subjects reported beliefs strategically or not. Remember the problem is that if subjects are strategic in reporting their beliefs they will report zero beliefs in an effort to increase their game payoffs (at the expense of their belief-elicitation payoffs). Our data suggest that this hardly ever happens. In fact, buyers never reported zero beliefs for message m_1 in any treatment, while sellers did so just a few times (less than 2%) in the Competition treatment and never in the No Competition treatment. This evidence strongly suggests that while the task of eliciting actions and beliefs in a psychological game is theoretically challenging, it was not something that dawned on our subjects nor should it since we arranged payments so as to create 'almost complete' incentive compatibility. This is, however, an issue that needs to be addressed when one tests psychological games in the lab by inducing psychological payoffs.

[29] To illustrate this approach, consider a buyer in the No Competition treatment, who believes that there is p_{m_i} chance that message m_i is sent by a high-quality seller and, instead, reports r_{m_i} in the experiment. The quadratic scoring rule we used in the experiment takes the form $c_1 - c_2 \cdot \text{mistake}^2$, where $c_1 = 100$ and $c_2 = 50$. Thus, a buyer's expected payoff in the belief task is

$$\mathbb{E}\Pi^{\text{beliefs}}\left(p_{m_i}, r_{m_i}\right) = p_{m_i} \cdot \left[c_1 - c_2\left((1-r_{m_i})^2 + (0-(1-r_{m_i}))^2\right)\right] + (1-p_{m_i}) \cdot \left[c_1 - c_2\left((0-r_{m_i})^2 + (1-(1-r_{m_i}))^2\right)\right] =$$
$$= p_{m_i} \cdot \left[c_1 - 2c_2(1-r_{m_i})^2\right] + (1-p_{m_i}) \cdot \left[c_1 - 2c_2(r_{m_i})^2\right].$$

The buyer's payoff from playing the game is

$$\mathbb{E}\Pi^{\text{game}}\left(p_{m_i}, r_{m_i}, \omega\right) = \begin{cases} 10p_{m_i} + (1-p_{m_i}) \cdot (-10\omega \cdot r_{m_i}) & \text{if this payoff is greater than 5} \\ 5 & \text{otherwise} \end{cases}$$

A risk-neutral buyer should report belief $r^*_{m_i}$ that maximizes his overall expected payoff

$$\mathbb{E}\Pi^{\text{Buyer}}\left(p_{m_i}, r_{m_i}, \omega\right) = \frac{1}{2} \cdot \mathbb{E}\Pi^{\text{belief}}\left(p_{m_i}, r_{m_i}\right) + \frac{1}{2} \cdot \mathbb{E}\Pi^{\text{game}}\left(p_{m_i}, r_{m_i}, \omega\right)$$

Given our parameters, the highest distortion in beliefs is $\max|p_{m_i} - r^*_{m_i}| = \frac{5}{2c_2} \cdot \left(1 - \frac{1}{\sqrt{2}}\right)$ is quite small and does not exceed 1.5%. Moreover, it results in a minimal increase in the buyer's payoff relative to reporting the true belief. In other words, our payment scheme is "practically" incentive compatible. We refer the reader to Section 4 in the Online Appendix, where we describe this procedure in detail.

4 Results

This section describes the performance of markets across our three treatments. We start by investigating the effects of induced psychological payoffs and competition on trade frequencies and market participants' payoffs. We then document buyers' and sellers' strategies and show which psychological types of sellers are mostly affected by sellers' competition. We conclude this section by comparing the outcomes in each market with those predicted by the theory to see if any of the equilibria organizes observed data in a satisfactory manner. Section 5 offers an explanation of the main forces driving the outcomes documented in this section.[30]

4.1 Trade

Figure 4 depicts the frequency of buyers' purchasing decisions in each treatment and the quality of the purchased goods. While all three treatments display similar outcomes in the first half of the experiment, once subjects have had the time to learn the game their behavior diverges and we document significant differences between the treatments.

For example, the presence of induced psychological payoffs significantly increases the trade frequency (Monetary vs No Competition, last 5 blocks: $p = 0.003$) and has a positive but not statistically significant effect on the average quality of sold goods (Monetary vs No Competition, last 5 blocks: $p = 0.119$). The additional introduction of competition between sellers further increases the trade frequency but lowers the quality of sold goods (No Competition vs Competition, last 5 blocks: $p < 0.001$ for trade frequency and $p = 0.098$ for quality of purchased goods). In fact, the competition between sellers undoes the benefits of induced psychological payoffs in the sense that it leaves the Buyers with goods that are comparable in quality to those purchased in the Monetary treatment (Monetary vs Competition, last 5 blocks: $p = 0.995$).

The combination of high trade frequency and low quality of purchased goods in markets with multiple sellers affects participants' payoffs. Table 2 reports regressions that document the effect of competition on buyers' and sellers' game payoffs in the two treatments with psychological payoffs. To make a comparison between the sellers' payoffs in the two treatments appropriate, we focus on the payoffs of the selected seller in the Competition treatment, because the non-selected seller earns a fixed payoff of zero. Also, in these regressions, we abstract away from the payoffs that subjects earn in the belief-elicitation task and focus only on the tree-game payoffs.

A few interesting patterns emerge from Table 2. Although no differences exist in the average payoffs of buyers in the first half of the experiment, buyers earn significantly less in the game with competition in the second half. In other words, buyers suffer from the presence of competition.

[30]Throughout this section, we use regression analysis to compare average outcomes between two groups (be that two treatments or two different types of Sellers). Specifically, we run random-effects GLS or LOGIT regressions (depending on the nature of the dependent variable) in which we regress the variable of interest (e.g., purchasing decision of buyers or the quality of the sold product) on a constant and a dummy variable that indicates one of the considered groups (i.e., two treatments or two messages), while clustering observations by sessions to account for potential interdependencies of observations within a session. We say that there is a significant difference between the two considered groups if the estimated coefficient on the dummy variable is significantly different from zero, and we report the p-value associated with it.

Most of the analysis presented below focuses on the last five blocks of each experimental session because subjects often learn the game by playing it. For this reason, the data from the first iterations of the game tend to be noisier because subjects are trying to figure out their strategies. By the second half of the experiment, subjects had experienced the game many times and may have possibly converged to their preferred strategies. However, we also present subjects' behavior in the first five blocks in several figures and tables to highlight changes in subject behavior.

Figure 4: Aggregate Outcomes, by treatment

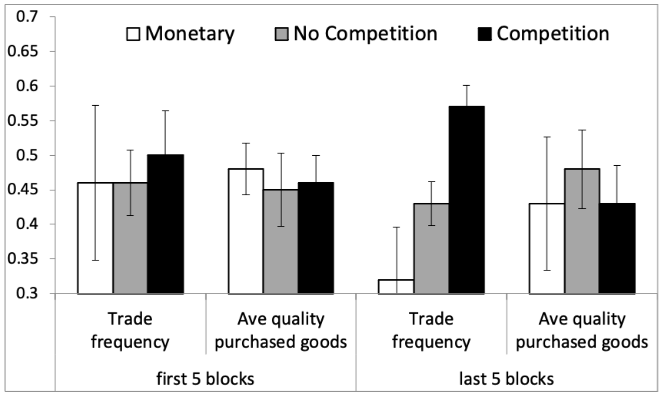

Notes: The left panel focuses on the first 5 blocks and the right panel on the last 5 blocks of the experiment. In each panel, we depict the trade frequency and the likelihood that the product was high quality conditional on the product being purchased. Bars indicate 95% confidence intervals using robust standard errors, which are computed by clustering observations by session.

Table 2: Effect of Competition on Payoffs of Buyers and (selected) Sellers in the Tree Game with Psychological Payoffs

	Buyers' Payoffs		Sellers' Payoffs	
	first 5 blocks	last 5 blocks	first 5 blocks	last 5 blocks
Competition treatment	-0.11 (0.17)	-0.77** (0.19)	-2.00** (0.60)	-2.44** (0.35)
Block number	-0.06 (0.06)	-0.05 (0.06)	-0.17 (0.12)	0.30** (0.12)
Constant	4.59** (0.21)	4.98** (0.52)	9.36** (0.56)	6.42** (1.02)
# of obs	2450	2450	2450	2450
# of clusters	7	7	7	7

Notes: Random-effects GLS regressions with the dependent variable being Buyers' payoffs in the tree game in the first two columns and Sellers' payoffs in the tree game in the last two columns. In all regressions, we abstract away from the payoffs that subjects accumulate for guessing beliefs tasks. Standard errors are clustered at the session level. ** indicates significance at the 5% level.

As for the sellers, we first note a clear ranking of average sellers' payoffs in the Competition treatment: selected sellers' payoffs are significantly higher than zero, which is what the non-selected sellers earn. Second, sellers suffer from the competition from the start of the experiment: in both halves of the experiment, selected sellers in the Competition treatment earn less than sellers in the No Competition treatment. These effects are statistically significant and large in magnitude and constitute the main punchline of our paper as summarized in Result 1 below.

Result 1: *Trade in markets without induced psychological payoffs is infrequent and when it happens most of the purchased goods are low-quality. The introduction of psychological payoffs increases* *both the trade and marginally increases the quality of purchased goods. The further introduction of competition between sellers results in even higher trade frequency but lowers the quality of purchased goods and, overall, negatively affects the payoffs of both buyers and sellers.*

4.2 Strategies of Buyers and Sellers

In this section, we look under the hood of our results and examine the strategies used by buyers and sellers. We do this by investigating the way sellers communicate with buyers, i.e., what messages sellers attach to goods of different qualities, the way buyers translate sellers' messages into purchasing decisions, and the resulting quality of purchases. Table 3 reports these statistics.

Table 3: Messages and Purchasing Decisions, by treatment

		Monetary	No Competition	Competition
first 5 blocks				
Sellers' communication behavior	$\Pr[m_1 \mid q = q_H]$	0.96 (0.02)	0.90 (0.06)	0.85 (0.02)
	$\Pr[m_1 \mid q = q_L]$	0.60 (0.07)	0.28 (0.03)	0.59 (0.02)
Buyers' purchasing behavior	$\Pr[\text{Buy} \mid m_1]$	0.56 (0.07)	0.59 (0.03)	0.57 (0.04)
	$\Pr[\text{Buy} \mid m_0]$	0.08 (0.02)	0.32 (0.04)	0.35 (0.03)
Quality of purchased goods	$\Pr[q = q_H \mid m_1 \& \text{Buy}]$	0.50 (0.02)	0.64 (0.02)	0.51 (0.02)
for different messages	$\Pr[q = q_H \mid m_0 \& \text{Buy}]$	n/a	0.08 (0.04)	0.26 (0.04)
last 5 blocks				
Sellers' communication behavior	$\Pr[m_1 \mid q = q_H]$	0.92 (0.01)	0.89 (0.06)	0.89 (0.01)
	$\Pr[m_1 \mid q = q_L]$	0.62 (0.02)	0.24 (0.03)	0.63 (0.03)
Buyers' purchasing behavior	$\Pr[\text{Buy} \mid m_1]$	0.39 (0.04)	0.56 (0.05)	0.65 (0.02)
	$\Pr[\text{Buy} \mid m_0]$	0.13 (0.06)	0.32 (0.04)	0.34 (0.03)
Quality of purchased goods	$\Pr[q = q_H \mid m_1 \& \text{Buy}]$	0.47 (0.02)	0.73 (0.02)	0.48 (0.03)
for different messages	$\Pr[q = q_H \mid m_0 \& \text{Buy}]$	0.06 (0.06)	0.06 (0.03)	0.15 (0.05)

Notes: The average observed quantities are presented with robust standard errors in parentheses. Standard errors are clustered at the session level.

Table 3 shows that there is little variation in sellers' communication strategies across our three treatments when subjects own a high-quality product; the vast majority of them send messages m_1.[31] It is the behavior of sellers with low-quality goods that differs. For example, the introduction of psychological payoffs reduces the lying frequency for low-quality sellers from 62% in the Monetary treatment to 24% in the No Competition treatment (last 5 blocks: $p < 0.001$). However, once the competition between sellers is introduced, the disciplining effect of psychological payoffs vanishes, and sellers with low-quality products lie as much in the Competition treatment as they do in the Monetary treatment (last 5 blocks: $p = 0.841$).

Turning to the buyers' behavior, we find that buyers' purchasing decisions after receiving mes-

[31] Focusing on the last 5 blocks of the experiment, the frequency of sending an m_1 message by a high-quality seller in any pair of treatments is not significantly different with $p = 0.892$ for Monetary vs No Competition treatment, $p = 0.095$ for Monetary vs Competition treatment, and $p = 0.510$ for No Competition vs Competition treatment.

sage m_1 are quite similar across the treatments in the first half of the experiment.[32] However, in the Monetary treatment, the buyers learn to purchase goods less after message m_1 as they gain experience (56% in the first 5 blocks vs 39% in the last 5 blocks, $p < 0.001$). There is no such declining trend in the markets with psychological payoffs. On the contrary, buyers in the Competition treatment learn to purchase goods with an m_1 label more often in the second half compared to the first half of the experiment ($p < 0.001$). As a result, in the second half of the experiment, buyers are less likely to purchase a product after message m_1 in the Monetary treatment compared with either No Competition or Competition treatments (Monetary vs No Competition: $p = 0.001$; Monetary vs Competition: $p < 0.001$). Furthermore, the purchasing frequencies for m_1 messages are slightly higher in the Competition than in the No Competition treatment in the second half of the experiment ($p = 0.018$) and are comparable for the m_0 messages ($p = 0.648$).[33]

Combining strategies used by two sides of the market, we compute the average quality of purchased goods that are advertised as 'really high-quality goods', i.e., those that come with the message m_1. Roughly 50% of these goods are actually high-quality goods in both the Monetary and the Competition treatment (last 5 blocks: $p = 0.829$). Contrary to that, in the No Competition treatment, most of the sold goods that come with message m_1 (73%) are actually high-quality goods (Monetary vs No Competition, last 5 blocks: $p < 0.001$ and Competition vs No Competition, last 5 blocks: $p < 0.001$).

Result 2: *The sellers with low-quality goods lie the least in the markets with induced psychological payoffs and no competition. The same types of sellers lie substantially more often in the other two market structures. At the same time, the buyers purchase more products with m_1 labels in markets with competition than in both types of markets without competition.*

4.3 The Effect of Psychological Types on Strategies and Payoffs

To understand the large difference between lying frequencies of low-quality sellers in markets with induced psychological payoffs with and without competition, we look into the sellers' communication strategy conditional on their psychological types. Figure 5 presents the average frequencies of m_1 messages sent by sellers of each type in the last five blocks of the experiment.[34]

Consistent with the averages presented in Table 3, we find that the vast majority of sellers who own a high-quality product disclose it truthfully and send message m_1. In fact, the 95% confidence intervals around the average frequencies of the m_1 message contain 100% for psychological types S3 and S4 in the No Competition treatment and for types S2, S3, and S4 in the Competition treatment. The deviations of the remaining types from being 100% truthful are rather small and are not very surprising given that any small tremble would be one-sided because of the boundary.

The situation changes when we look at sellers with low-quality goods. In the No Competition treatment, we find that about 50% of low-quality sellers with type S1 choose to lie, whereas other types (S2, S3, and S4) lie much less. By contrast, in the Competition treatment, both types S1 and S2 of the sellers with low-quality products lie the majority of the time (about 80% of types S1 and

[32]In the first half of the experiment, we detect no significant difference in the likelihood of buying a product conditional on receiving the m_1 message between any pair of treatments, $p > 0.10$ in all three pair-wise comparisons.

[33]In the second half of the experiment, the buyers purchase the product with the m_0 message more often in the two markets with induced psychological payoffs than in the market with only monetary payoffs (Monetary vs No Competition (Competition), last 5 blocks: $p = 0.046$ ($p = 0.014$).

[34]The sellers' communication decisions in the first five blocks look very similar (see Figure 5 in the Online Appendix).

Figure 5: Communication Decisions of Sellers in Markets with Psychological Payoffs, last 5 blocks

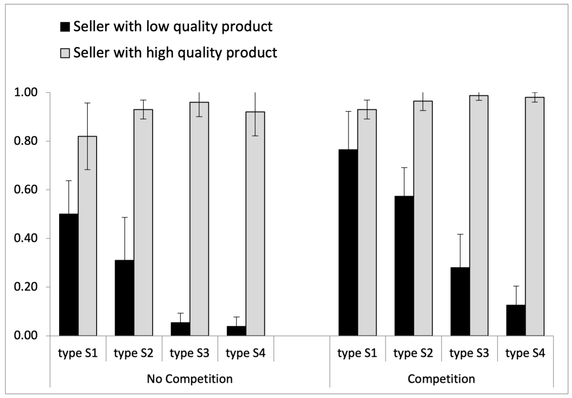

Notes: Average frequency of sending message m_1 is presented for each type of Seller in each treatment in the second half of the experiment. We compute 95% confidence intervals using robust standard errors obtained by clustering observations by session.

about 60% of types S2), whereas types S3 and S4 lie much less.[35] In fact, in both treatments, we observe a monotonic decrease in the lying frequency of low-quality sellers as we move from type S1 to type S4.[36]

More importantly, sellers change their behavior when competition is introduced. Sellers with low-quality goods lie significantly *more* in the Competition than in the No Competition treatment for all four possible psychological types they may have. For example, in the No Competition treatment, sellers with types S1, S2, S3, and S4 sent untruthful m_1 messages conditional on having a low-quality good, 50%, 31%, 5%, and 4% of the time, respectively, while these percentages increased to 77%, 57%, 28%, and 13%, when the competition was present. Pairwise comparisons between these fractions confirm the directional results evident in Figure 5 ($p = 0.012$ for types S1, $p = 0.005$ for types S2, $p < 0.001$ for types S3, and $p = 0.044$ for types S4).

Frequent lying of sellers with high psychological costs in the Competition treatment affects their overall payoffs in the game. Table 4 addresses this point and shows which types of sellers and buyers suffer the most from the presence of competition looking at the last 5 blocks of the experiment.[37]

As Table 4 shows sellers who own the high-quality product earn the same average payoffs in both treatments, irrespective of their psychological type. The sellers with low-quality goods

[35] The fact that, in the No Competition treatment, the S1 sellers with the low-quality products are truthful about half of the time is consistent with the idea that some participants experience an additional home-grown aversion to lying that they brought to the experiment since S1 is the type for which we did not induce psychological payoffs ($L = G = 0$). Interestingly, the competition between the sellers dominates this home-grown aversion to lying since the same types lie about 80% of the time in the markets with multiple sellers.

[36] In the last five blocks of the No Competition treatment, low-quality sellers with S1 types lie significantly more than S2 types ($p = 0.003$), S2 types lie significantly more than S3 types ($p < 0.001$), whereas there is no significant difference between lying frequencies of S3 and S4 types ($p = 0.556$). In the last five blocks of the Competition treatment, low-quality sellers with S1 types lie significantly more than S2 types ($p < 0.001$), S2 types lie significantly more than S3 types ($p < 0.001$), and S3 types lie significantly more than S4 types ($p < 0.001$).

[37] Table 1 in the Online Appendix presents the same statistics for the first 5 blocks of the experiment.

Table 4: Which Types of Buyers and Sellers Suffer the Most from Competition (last 5 blocks)?

		No Competition	Competition	Difference
	type S1 ($G=0, L=0$)	12.20 (0.84)	14.93 (1.03)	YES** ($p=0.05$)
SELLERS	type S2 ($G=6, L=0$)	8.26 (0.54)	6.67 (0.76)	YES* ($p=0.08$)
low-quality product	type S3 ($G=0, L=20$)	9.61 (0.65)	0.08 (1.11)	YES** ($p<0.01$)
	type S4 ($G=6, L=20$)	8.07 (1.22)	-0.70 (1.94)	YES** ($p<0.01$)
	type S1 ($G=0, L=0$)	7.46 (0.28)	7.83 (0.23)	NO ($p=0.20$)
SELLERS	type S2 ($G=6, L=0$)	7.69 (0.26)	8.52 (0.41)	NO ($p=0.15$)
high-quality product	type S3 ($G=0, L=20$)	7.90 (0.27)	7.99 (0.23)	NO ($p=0.72$)
	type S4 ($G=6, L=20$)	7.56 (0.23)	7.77 (0.26)	NO ($p=0.54$)
	$\omega \leq 0.2$	4.15 (0.33)	4.14 (0.47)	NO ($p=0.96$)
	$0.2<\omega \leq 0.4$	4.11 (0.33)	3.37 (0.43)	NO ($p=0.18$)
BUYERS	$0.4<\omega \leq 0.6$	4.74 (0.25)	4.08 (0.43)	NO ($p=0.17$)
	$0.6<\omega \leq 0.8$	4.78 (0.22)	2.97 (0.35)	YES** ($p<0.01$)
	$\omega>0.8$	5.04 (0.13)	4.54 (0.21)	YES** ($p=0.04$)

Notes: We report average payoffs of buyers and sellers in the last five blocks of the experiment and the robust standard error in parentheses. The last column reports the results of a statistical test comparing payoffs for a fixed type of buyer or seller in the two treatments. * and ** indicate significance at the 10% and the 5% levels, respectively.

are the ones who suffer from the competition because they earn significantly lower payoffs for all four psychological types. The largest losses are experienced by sellers with low-quality goods and psychological types S3 and S4 who have a strong aversion to lying and might have a strong sensitivity to guilt. Interestingly, sellers with low-quality products and psychological types S3 or S4 who are selected by buyers to play the tree game earn average payoffs that are not statistically different from zero, which in fact is the outside option for a seller not engaging in trade. As for the buyers, those with higher disappointment aversion suffer the most losses from competition between sellers.

Result 3: *In markets with induced psychological payoffs, sellers with higher values of guilt and lying sensitivity lie less than those with lower values. Competition leads to more lying by sellers who own low-quality products, irrespective of their psychological type. Furthermore, competition negatively affects the payoffs of buyers with high sensitivity to disappointment and payoffs of sellers with low-quality goods and high lying or guilt aversion.*

4.4 Equilibrium predictions

We finish this section by comparing the behavior of our market participants to the equilibrium predictions (Section 2.4). We start with the pooling equilibrium, which predicts that sellers' messages are uncorrelated with the product quality, buyers treat messages as uninformative and ignore them when making purchasing decisions, and, as a result, no trade should occur. Given that all the characteristics of the pooling equilibrium are extreme, it is not surprising that neither of our treatments closely tracks these predictions. However, we find that the Monetary treatment is the closest among all three treatments to the above predictions. Indeed, in the last 5 blocks of the experiment, the majority of both buyers and sellers in markets without induced psychological payoffs play the pooling equilibrium strategies: 61% of the sellers and 55% of the buyers.[38]

The fact that some of our results in the Monetary treatment deviate from what is expected in a perfect pooling equilibrium points out that, although we were able to make lying and guilt more salient in our games with induced psychological payoffs, imposing no penalties for lying or guilt does not mean those forces do not exist. Some subjects simply do not like to lie, and they feel guilty if they do so, which is why a model that assumes people are capable of abstracting away from such psychological costs may not be realistic.

The behavior of participants in the markets with psychological payoffs clearly rejects the hypothesis that they are playing the pooling equilibrium. Indeed, sellers' communication strategies are informative of the product quality (Figure 5) and buyers incorporate this information when making their purchasing decisions (Table 3). Individual level analysis confirms the same point: less than 10% of the sellers and less than 30% of the buyers in either treatment play the pooling equilibrium strategies (see Tables 2 and 3 in the Online Appendix).

Having dispensed with the pooling equilibrium as the one compatible with the No Competition and Competition data, we turn next to the two partially informative equilibria: the PIE1, in which two of the four psychological types of sellers with low-quality products lie in equilibrium, and the PIE2, in which only one type does, the one that experiences no guilt and no lying aversion. We ask whether one of these two equilibria organizes data in a satisfactory manner. To do that, we present in Table 5 a subset of PIE1 and PIE2 predictions and compare them to the observed outcomes. This subset contains those predictions that differ across these two equilibria.

Table 5: Aggregate Fit of Equilibrium Predictions, last 5 blocks

	Observed	Predicted		Statistical Tests	
		PIE1	PIE2	Observed vs PIE1	Observed vs PIE2
No Competition treatment					
Pr[trade]	0.43 (0.02)	0.36	0.55	$p < 0.01$	$p < 0.01$
$\Pr[m_1 \mid q = q_L]$	0.24 (0.03)	0.50	0.25	$p < 0.01$	$p = 0.62$
$\Pr[\text{Buy} \mid m_1]$	0.56 (0.05)	0.51	1.00	$p = 0.32$	$p < 0.01$
$\Pr[q = q_H \mid m_1 \&\ \text{trade}]$	0.73 (0.02)	0.57	0.73	$p < 0.01$	$p = 0.88$
Competition treatment					
Pr[trade]	0.57 (0.02)	0.46	0.80	$p < 0.01$	$p < 0.01$
$\Pr[m_1 \mid q = q_L]$	0.63 (0.03)	0.50	0.25	$p < 0.01$	$p < 0.01$
$\Pr[\text{Buy} \mid m_1]$	0.65 (0.02)	0.51	1.00	$p < 0.01$	$p < 0.01$
$\Pr[q = q_H \mid m_1 \&\ \text{trade}]$	0.48 (0.03)	0.57	0.73	$p = 0.01$	$p < 0.01$

Notes: The first column reports observed frequencies of trade, m_1 messages, and purchasing decisions of buyers focusing on those that are different across the two partially informative equilibria. The second and third columns report theoretically predicted frequencies. The last two columns report the p-values comparing observed frequencies with those predicted by PIE1 and PIE2.

The No Competition treatment conforms to some PIE2 predictions but not to all. In particular,

[38] A pooling-equilibrium strategy for a seller means choosing the same message irrespective of the product quality and one's psychological type if such is induced. A pooling-equilibrium strategy for a buyer means making the same purchasing decision irrespective of the received message in the Monetary treatment and choosing two "similar" purchasing cutoffs in the No Competition and Competition treatments, where similar means cutoffs are no more than 0.05 away from each other.

the sellers' behavior and the average quality of purchased goods that come with the m_1 label match theoretically predicted levels. However, buyers purchase goods significantly less than what PIE2 predicts. This leads to a lower frequency of overall trade. At the same time, the data in the Competition treatment is clearly incompatible with theoretical predictions of either the PIE1 or PIE2 equilibrium. Hence, we must conclude the difference in behavior between the Competition and the No Competition treatment can not be attributed to subjects selecting different equilibria in these different treatments.

Result 4: *Majority of buyers and sellers play the pooling equilibrium strategies in markets without induced psychological payoffs. The markets with induced psychological payoffs and no competition conform to predictions of the most informative equilibrium (PIE2) with the exception of lower-than-predicted trade frequency. The markets with induced psychological payoffs and competition between sellers do not match either of the equilibria.*

5 Main forces driving markets to these outcomes

Our results so far have generated a number of puzzles that need to be discussed. The first is why buyers in markets with psychological costs but no competition are hesitant to buy goods after receiving an m_1 message (i.e. why do they set such low $\omega's$) yet the behavior of sellers in this treatment is basically consistent with the most informative equilibrium (i.e., is relatively honest). Second, why competition between sellers leads to more lying by sellers? Third, why do buyers purchase products with the label m_1 more often in markets with competition when sellers in those markets are more prone to lying?

In this section we attempt to offer some solutions to these puzzles but before we do let us pause and discuss the beliefs of our seller and buyer subjects across our Competition and No Competition treatments. These beliefs will be a key ingredient for describing the subject's behavior.

Beliefs. Beliefs are an important part of explaining the puzzles mentioned above since behavior is supposed to be the best response to them. Recall that in addition to strategies, we elicit buyers' first-order beliefs of what messages mean and sellers' second-order beliefs regarding buyers' first-order beliefs. Specifically, buyers indicate the likelihood that message m_i was sent by a high-quality seller, while the sellers try to guess the likelihood stated by the buyers, i.e., buyers' interpretation of messages. Table 6 presents summary statistics of elicited beliefs and compares those to actual meanings of messages in the last 5 blocks of the experiment, while Table 4 in the Online Appendix replicates the same analysis for the first 5 blocks.

One interesting fact that stands out in Table 6 is that while buyers correctly predict the lying behavior of sellers in markets without competition, they grossly overestimate the honesty of sellers when competition is introduced. Obviously, if competition leads buyers to be overly optimisitc in their beliefs about sellers, it is not surprising, as we will see, that they ultimately will purchase goods when they shouldn't. More precisely, in markets without competition, the average beliefs for m_1 message are 0.76 which are not statistically different from the veracity of such messages which is 0.71 ($p > 0.05$). Moreover, average buyers' beliefs match those predicted by the PIE2, which organizes the data in the No Competition treatment quite well.[39] However, in markets with competition between sellers, buyers do quite poorly at predicting the average quality of the product

[39] According to PIE2, 73% of m_1 messages come from a high-quality seller. We cannot reject the null that the observed buyers' beliefs are significantly different from this number with $p = 0.259$.

Table 6: Buyers' and Sellers' Beliefs, Buyers' Purchasing Cutoffs, and Actual Quality of Products for Different Messages, last 5 blocks

	$\omega'(m_i)$	$\bar{b}^1_B(m_i)$	$\bar{b}^2_S(m_i)$	$\bar{q}_H(m_i)$	$\bar{b}^1_B(m_i)$ = $\bar{b}^2_S(m_i)$	$\bar{b}^1_B(m_i)$ = $\bar{q}_H(m_i)$	$\bar{b}^2_S(m_i)$ = $\bar{q}_H(m_i)$
No Competition							
message m_0	0.29 (0.03)	0.26 (0.06)	0.20 (0.04)	0.07 (0.04)	$p = 0.009$	$p = 0.001$	$p = 0.001$
message m_1	0.59 (0.05)	0.76 (0.03)	0.73 (0.02)	0.71 (0.02)	$p = 0.136$	$p = 0.068$	$p = 0.247$
Competition							
message m_0	0.31 (0.03)	0.22 (0.03)	0.25 (0.02)	0.17 (0.02)	$p = 0.318$	$p = 0.167$	$p < 0.001$
message m_1	0.62 (0.04)	0.77 (0.03)	0.70 (0.03)	0.49 (0.03)	$p = 0.004$	$p < 0.001$	$p < 0.001$

Notes: The first column records average cutoffs reported by buyers for each message, which is the highest disappointment sensitivity for which a buyer is willing to purchase the product that comes with message m_i. The second and third columns, $\bar{b}^1_B(m_i)$ and $\bar{b}^2_S(m_i)$, are buyers' first-order and the sellers' second-order beliefs for message m_i. The fourth column, $\bar{q}_H(m_i)$, is the likelihood that message m_i comes from the high-quality seller estimated using the actual realizations observed in each round of each block. In all cells, the robust standard errors are reported in parentheses. The last three columns report results of statistical tests comparing buyers' and sellers' beliefs (fifth column), buyers' beliefs and the average actual frequency of high-quality sellers for different messages (sixth column), and sellers' beliefs and the average actual frequency of high-quality sellers for different messages (seventh column).

that comes with label m_1: average buyers' beliefs are almost 30 percentage points above the actual number (0.77 versus 0.49). In other words, buyers significantly overestimate how good products are when they come with the message m_1 in markets with multiple sellers.

While sellers are never supposed to lie when sending the m_0 message (such beliefs are expected to be zero since only a low-quality seller is supposed to use message m_0), we find that in the experiment such beliefs are 0.26 and 0.22 in the No Competition and Competition treatment in the last 5 blocks, respectively. Given that zero is a corner solution, unsurprisingly, we find that all the deviations are positive and move average buyers' beliefs away from the prediction of zero. A perhaps more informative statistic might be a fraction of times that reported beliefs were close to zero, allowing for some small noise. In the last five blocks of the No Competition treatment, the majority (63% in the No Competition and 77% in the Competition treatment) of reported beliefs upon observing message m_0 are at most 5 percentage points away from zero.

Despite the difference in seller behavior across markets with and without competition, it is interesting that buyer beliefs are remarkably similar which suggests that buyers did not adjust their beliefs to accommodate the different environment they were in. This can be seen by focusing on the last five blocks of the experiment and noting that after observing an m_0 message, the mean belief of buyers was 0.26 in the No Competition treatment and 0.22 in the Competition treatment. After receiving the m_1 message, these beliefs were 0.76 and 0.77, respectively. These beliefs are statistically indistinguishable across the two treatments ($p = 0.462$ for message m_0 and $p = 0.927$ for message m_1).[40]

[40] The similarity in buyers' beliefs in the two treatments is consistent with the similarity in buyers' purchasing decisions. Conditional on receiving an m_0 message, buyers choose cutoffs of 0.29 and 0.31 for the No Competition and Competition treatment in the last 5 blocks, which are not statistically different from each other ($p = 0.627$). Upon receiving the m_1 message, buyers' cutoffs in the last 5 blocks were 0.59 and 0.62 for the No Competition and Competition treatments, which are also not different ($p = 0.558$).

The performance of our markets also depends on the beliefs of sellers. These are second-order beliefs which are the beliefs of sellers about the first-order beliefs of buyers. In looking at these beliefs, we observe that sellers correctly predict buyers' beliefs for the m_1 message in the No Competition treatment ($p = 0.136$) and for the message m_0 in the Competition treatment ($p = 0.318$).[41] For the remaining two cases, message m_0 in the No Competition treatment and message m_1 in the Competition treatment, sellers think that buyers' beliefs are slightly lower than they actually are, but the differences are quite small in magnitude. Combining data for the No Competition treatment presented in Tables 5 and 6 we find that sellers' overall frequency of lying when they own a low-quality product is consistent with the beliefs they hold regarding buyers' interpretation of the m_1 messages.

With this background let us proceed to discuss our three puzzles.

5.1 Puzzle 1: Buyers' Beliefs and Purchasing Decisions in Markets without Competition

Why do buyers in markets with no competition hesitate to buy goods after receiving the m_1 message, especially when those messages are basically honest and they believe them to be so?[42]

A simple explanation can be found if we allow buyers to be risk-averse. In this case, the optimal cutoffs for a such buyer with "correct" beliefs should be lower than for the risk-neutral buyers we assumed in our model and hence lead them to be more cautious in their buying behavior.

To test whether our data support this explanation, we use an additional measure of risk aversion (an investment task), which we collected at the end of each experimental session. In the investment task, subjects were asked to allocate a budget of 200 points between a risk-free asset that paid one point for every point invested and a risky asset that paid 2.5 or 3 points with a probability of 0.50 or 0.40 for each point invested in the investment task 1 and 2, respectively. Thus, subjects who invest the full amount in the risky asset are either risk neutral or risk-loving, whereas lower than full investment indicates a subject is risk-averse. In addition, we can rank subjects in terms of their risk attitudes: the lower the amount invested in the risky asset, the more risk averse she is. Our data indicate a significant correlation between buyers' purchasing cutoffs upon observing m_1 and their risk attitudes: buyers who are more risk averse set lower cutoffs for m_1 (No Competition treatment, last 5 blocks: $p = 0.026$).[43]

[41] The fact that in the No Competition treatment sellers' beliefs for m_1 match average buyers' beliefs and also match beliefs predicted by the PIE2 adds one more piece of evidence that the most informative equilibrium, PIE2, predicts the average behavior in the markets without competition quite well.

[42] The second row in Table 6 indicates that the average cutoff for m_1 message is 0.59. This is consistent with the numbers reported in Table 3 that shows that given actual realizations of types in the experiment, buyers purchased 56% of goods that came with message m_1. The PIE2 predicts that buyers should always purchase the product that comes with the message m_1.

[43] To reach this conclusion, we use the ORIV (Obviously Related Instrumental Variables) technique developed by Gillen et al. (2019), which allows us to correct for the measurement errors in the elicitation of risk attitudes. ORIV is an improved version of the traditional instrumental variables approach to errors-in-variables, which produces consistent coefficients, correlations, and standard errors and an estimator that is more efficient than standard instrumental variable techniques. The ORIV estimation was performed for the average m_1 cutoff stated by buyers in the last 5 blocks of the No Competition treatment.

5.2 Puzzle 2: Seller Behavior in Markets with Competition

Our second puzzle is why do sellers lie more often in the markets with competition than in the markets without competition given that buyers behave the same way in the two markets and sellers are well aware of it?

Note that in markets with multiple sellers, the optimal actions of sellers depend not only on their beliefs about buyers' purchasing behavior but also on their beliefs about the other sellers' actions. Such considerations are not present in markets with one seller and introduce a layer of strategic uncertainty that requires each seller to form beliefs about the other seller's actions.

It is our contention that sellers lie more often in markets with competition because they feel they have to in order to sell their goods. This is, in part, a result of the feedback they get during their interaction with the market. For example, sellers observe all messages sent by their competitors but they only observe the quality of competitors' goods when the purchase happens. Thus, a seller does not always know whether the other seller sent a message m_1 because he truly owns a high-quality product or he is lying. What sellers do know from experience is that buyers tend to select a seller with the m_1 message when facing two different messages, and, as a result, the excluded seller gets a zero payoff.[44]

Table 7 shows evidence consistent with the logic above. For each psychological type of a low-quality seller, we regress an indicator that such a seller lies to the buyer and sends a message m_1 in a particular block of the experiment. The right-hand side variables include two variables that capture the behavior of the competitor seller in the previous block and two variables that capture the behavior of the buyer. We control for the seller's own tendency to lie for the same psychological type in the previous block, the sellers' second-order beliefs, and the block number. The regression analysis shows that sellers react to competitive pressure by lying more often when their competitors send more m_1 messages irrespective of their psychological type and this effect is large in magnitude (the first row). In addition, low-quality sellers with low psychological costs, i.e., the S1 and the S2 types, lie more often in response to observing that their competitor unloads the low-quality good by sending the m_1 message (the second row).

As Table 7 shows, low-quality sellers with type S2, i.e., the type who suffers only from guilt aversion, react the most to the actions of their competitors. There is a theoretically grounded reason for why this might be the case which is that the presence of strategic uncertainty regarding the other seller's behavior pushes the S2 low-quality type seller towards lying more frequently.

To explains, consider a seller, seller A, who has some doubts about the other seller's strategy, i.e., seller B. In particular, seller A believes that seller B always sends the message m_1 if he owns a high-quality product and if he owns a low-quality product and has psychological type S1. Seller A also believes that seller B sends the message m_0 when he owns a low-quality product and has type S3 or S4. However, seller A is uncertain of what S2 low-quality seller B does and assigns probability x to seller B lying in this situation. In other words, seller A believes that seller B plays the PIE1 strategy with probability x and the PIE2 strategy with probability $1 - x$. Furthermore, assume that seller A holds empirically correct average beliefs about buyers' interpretation of the message m_1, i.e., $b^2_{\text{seller A}}(m_1) = 0.70$. The question is then what is the optimal response of seller A?

The simple calculation shows that seller A with a high-quality product or low-quality product and an S1 type both prefer to send the m_1 message. Moreover, the S3 and S4 types with low-quality

[44] When faced with two identical messages, buyers are equally likely to select either one of the two sellers, i.e., the probability of selecting the first seller is not significantly different from 50% (last 5 blocks: $p = 0.53$). However, when a buyer receives two different messages, then she selects the seller with m_1 message in 84% cases in the last 5 blocks; this frequency is significantly different from 50% ($p < 0.001$).

Table 7: Learning by Observing Actions of Other Sellers in the Competition treatment

	Dependent Variable: Indicator that low-quality seller with specific psych type sends m_1			
	S1 type	S2 type	S3 type	S4 type
Behavior of competitors				
frac of m_1 messages in the prev block	0.47** (0.09)	0.56** (0.11)	0.40** (0.10)	0.26** (0.08)
frac of lowQ products conditional on m_1 and buyer purchasing it in the prev block	0.14** (0.06)	0.21** (0.07)	0.07 (0.06)	0.04 (0.05)
Behavior of buyers				
how often m_1 is selected when diff messages	0.06 (0.05)	-0.01 (0.06)	0.16** (0.05)	0.08* (0.04)
how often buy when selected message is m_1	-0.01 (0.008)	-0.01 (0.01)	-0.02** (0.01)	-0.001 (0.01)
Nb obs	414	414	414	414
Nb subjects	46	46	46	46
Overall R-square	0.2156	0.2108	0.1400	0.1489

Notes: Random-effects GLS regressions with fixed effects for sessions and with robust standard errors clustered at the individual level. In all regressions, we control for the block number, for sellers' second-order beliefs in the current block, and for sellers' own behavior given the same type in the previous block. ** indicates significance at the 5% level.

products prefer to be truthful and send the m_0 message. However, even for small values of x, such a low-quality seller A with an S2 type would want to send an m_1 message, i.e., play PIE1. In other words, the best response of an S2 seller with a low-quality product who is unsure of which PIE his opponent is playing is to lie. This points to the fragility of PIE2 relative to PIE1 in the face of strategic uncertainty, which is surely present in both actual and experimental markets with multiple sellers.

5.3 Puzzle 3: Buyer Behavior in Markets with Competition

Our final puzzle is why, if sellers are lying more in markets with competition, the buyers in those markets do not adjust their beliefs accordingly. This is puzzling since in markets without competition we find that buyers do learn how to interpret messages correctly. We think there are several explanations for the sluggishness of beliefs in our Competition treatment. While some of these explanations are cultural, others are more rationality based and follow from simple belief updating.

To start with the simplest explanation, people may feel that competition is there to protect them and hence may start out the experiment trusting m_1 messages and resist altering these exaggerated beliefs when contrary evidence is observed, i.e these initial beliefs adjust sluggishly. Such a belief in competitive markets is widespread and inculcated into us from the day we are born.[45] Consistent with this idea, we observe that buyers' beliefs in the first block of the experiment are marginally higher in the Competition compared with the No Competition treatment, $b_B^1(m_1) = 0.76$ vs $b_B^1(m_1) = 0.71$, even though the two are not statistically different ($p = 0.508$). These first block beliefs represent buyers' initial instincts about sellers' behavior after reading the instructions and before accumulating any experience of playing the game.[46]

[45] The idea that competition is beneficial for consumers is common wisdom in democratic societies. The channels through which competition works to improve consumers' welfare include encouraging businesses and firms to reduce prices, improve the quality of existing goods, and stimulate the creation of new products. For example, see the European Commission report on why competition policies are important for consumers at `https://competition-policy.ec.europa.eu/consumers/why-competition-policy-important-consumers_en`.

[46] Notably, the average first beliefs in the No Competition treatment are statistically indistinguishable from the

The sluggish belief updating, however, is also consistent with rational behavior since in markets with competition buyers receive several messages and compare them before making a choice. In our experiment, for example, a laboratory buyer may receive two m_1 messages and choose one seller at random. If the good purchased turns out to be of low quality the buyer may not change his posterior too much about the truthfulness of m_1 messages in the market in general because the other seller, whose offer was rejected, may have been reporting truthfully and hence the downgrading of the truthfulness of the m_1 message in competitive markets may not be as drastic as it is in non-competitive markets where there is only one message which, if deceptive, would be a clear signal of the truthfulness of the m_1 message. Hence, it may be rational to update conservatively in competitive markets. We formalize this argument in a simple belief updating learning model for buyers.

To capture this logic more formally, consider a buyer who holds beliefs about sellers' strategies and updates them based on the information observed in the experiment in a Bayesian manner. In particular, assume a buyer believes that any seller with a high-quality good sends message m_1, type S1 with a low-quality good also sends m_1, since she experiences no psychological dis-utilities of any kind, and types S3 and S4 with low-quality goods send the truthful message m_0 because lying is too costly for them. That leaves the seller of type S2 as the only seller whose behavior is uncertain and it is that behavior that a buyer updates given an observed product's quality and message that came with it.

In a market with only one seller, a buyer believes such a low-quality seller with S2 type sends message m_0 with probability μ. A buyer's belief about a seller's strategy determines the main quantity of interest: the informativeness of m_1 message $b_B^1(m_1) = \Pr[q = q_H|m_1]$, which determines optimal purchasing behavior. This quantity can be written as

$$b_B^1(m_1) = \Pr[q = q_H|m_1] = \mu \cdot \frac{0.4}{0.4 + 0.6 \cdot 0.25} + (1-\mu) \cdot \frac{0.4}{0.4 + 0.6 \cdot 0.5} = \frac{44 + 12\mu}{77}$$

In a market with two sellers, a buyer allows for the possibility that some sellers might be more truthful than others. Specifically, a buyer believes that there is a 50% chance that a low-quality seller with type S2 sends m_0 with probability μ_1 and a 50% chance that this probability is $\mu_2 < \mu_1$. In this case,

$$b_B^1(m_1) = \Pr[q = q_H|m_1] = \frac{1}{2}\left(\frac{44 + 12\mu_1}{77} + \frac{44 + 12\mu_2}{77}\right)$$

Consider now what happens when a buyer observes an m_1 message that came from a low-quality seller. In a market without competition, a buyer updates downward his belief about μ, let's call this posterior belief μ', and uses this posterior to adjust the belief about the m_1 message in the future. Bayes' rule dictates $\mu' = \frac{\mu}{2-\mu} < \mu$, which implies that the posterior informativeness of m_1 message is lower than what the buyer originally thought it is. The overall change in buyer's belief as a result of observing an m_1 message coming from the low-quality seller is

$$\Delta_{\text{No Comp}} = \frac{44 + 12\mu}{77} - \frac{44 + 12\mu'}{77} = \frac{12}{77}(\mu - \mu') = \frac{12}{77} \cdot \frac{\mu(1-\mu)}{2-\mu}$$

In a market with competition, the buyer updates both μ_1 and μ_2 beliefs, and arrives at the lower posterior belief about an m_1 message as well. This new posterior is $Pr[q_H|m_1] = \frac{1}{2} \cdot \left(\frac{44+12\mu_1'}{77} + \frac{44+12\mu_2'}{77}\right)$ where $\mu_i' = \frac{\mu_i}{2-\mu_i}$ for $i = 1, 2$. Note, however, that if both markets start with

beliefs predicted by PIE2 ($p = 0.765$), which means buyers start with PIE2 equilibrium beliefs in this treatment.

the same prior belief about an m_1 message[47], which is the case in our data, then beliefs changes are *smaller* in the markets with than without competition:

$$\Delta_{\text{Comp}} = \frac{1}{2} \cdot \frac{12}{77} \cdot \left(\frac{\mu_1(1-\mu_1)}{2-\mu_1} + \frac{\mu_2(1-\mu_2)}{2-\mu_2} \right) < \Delta_{\text{No Comp}}$$

In other words, Bayesian updating dictates slower responsiveness in beliefs' in markets with than without competition.

Empirical evidence of beliefs' sluggishness. Given our expectation that buyer beliefs are updated more slowly in markets with competition, let us turn to our data for confirmation.

Table 8 presents regressions depicting how buyers adjust their beliefs regarding the m_1 message conditional on the feedback they observe in the experiment. For each treatment, the dependent variable is buyers' beliefs regarding m_1 in a particular block. The right-hand side variables include buyers' initial beliefs about m_1, their previous block beliefs about m_1, and the likelihood of purchasing a high-quality good conditional on receiving each possible message in the previous block.

Table 8: The Evolution of Buyers' Beliefs

	Beliefs in block t	
	No Competition	Competition
Beliefs		
$b_B^1(m_1)$ in block 1	0.06 (0.06)	0.35** (0.07)
$b_B^1(m_1)$ in block $t-1$	0.57** (0.06)	0.27** (0.07)
Feedback		
$\Pr[q = q_H \| m_1]$ in block $t-1$	0.08** (0.04)	0.09** (0.04)
$\Pr[q = q_H \| m_0]$ in block $t-1$	0.02 (0.07)	0.04 (0.04)
Nb obs	230	207
Nb participants	26	23
Overall R-sq	0.4303	0.3956

Notes: Random-effects GLS regressions with robust standard errors clustered at the individual level. $\Pr[q = q_H|m_i]$ is the fraction of high-quality products that came with the m_i label for $i = 1, 2$. To account for inter-dependencies of observations that come from the same session, we include session-fixed effects. Controls include block number, risk attitude measures, and beliefs about m_0 in the previous block.

Table 8 shows that in both treatments interpretations of m_1 messages depend on how often buyers observe high-quality products with the m_1 label in the previous block. This experience accumulates over time and affects how much buyers trust the future m_1 messages. However, in markets with competition, beliefs are more rigid and reflect to the large extent the very first belief that buyers formed before observing any sellers' behavior. This can be seen by the positive, significant, and large in magnitude estimated coefficient on $b_B^1(m_1)$ in the first block in the markets with competition regression (second column), while the same coefficient is not significant in the markets without competition (first column).

[47] In other words,

$$\frac{44+12\mu}{77} = \frac{1}{2} \cdot \left(\frac{44+12\mu_1}{77} + \frac{44+12\mu_2}{77} \right) \Leftrightarrow \mu = \frac{\mu_1}{2} + \frac{\mu_2}{2}$$

The buyers' interpretation of messages affects their purchasing decisions, and this is where the sluggishness of beliefs plays an important role. Table 9 shows the connection between beliefs and actions in the two markets. To do that we look at individual purchasing cutoffs for an m_1 message averaged over the last 5 blocks of the experiment, and ask how these average cutoffs depend on the changes in beliefs about the content of an m_1 message between the two halves of the experiment, the difference in interpretations of the m_1 and the m_0 messages, and the probability of choosing a seller with an m_1 over an m_0 message when multiple sellers are present in the market.

Table 9: Buyers' Purchasing Decisions

Dependent Variable: Purchasing cutoff for an m_1, last 5 blocks		
	No Competition	Competition
$b_B^1(m_1)\|_{\text{last 5}} - b_B^1(m_1)\|_{\text{first 5}}$	0.66** (0.13)	1.23* (0.45)
$b_B^1(m_1)\|_{\text{first 5}} - b_B^1(m_0)\|_{\text{first 5}}$	0.27 (0.27)	0.56* (0.22)
Prob choose m_1 over $m_0\|_{\text{last 5}}$		0.92** (0.09)
Nb obs	26	23
R-sq	0.3822	0.6528

Notes: Linear regressions with robust standard errors clustered at the session level. The right-hand side variable is the difference in average beliefs for m_1 message between the last 5 and the first 5 blocks of the experiment. The second one is the difference in average beliefs between m_1 and m_0 messages in the first 5 blocks. The last right-hand side variable is the chance that the buyer in the Competition treatment chooses a seller with the m_1 message over the m_0 message if two messages are different, averaged over the last 5 blocks. We control for the difference in beliefs between the two messages in the last 5 blocks of the experiment and the risk attitude of buyers.

In the markets without competition (first column), buyers' purchasing decisions are strongly affected by the change in beliefs regarding an m_1 message, which reflects their actual experiences in the markets and not just the initial beliefs. At the same time, in the markets with competition (second column), both actual experiences and initial beliefs play a role but these links are weaker as can be seen by a drop in the significance of both variables; both effects are marginal at the 10% level. The main determinant of buyers' purchasing cutoffs in these markets is how likely a buyer is to choose a seller with an m_1 message over the one with an m_0. In a sense, this last variable captures a buyer's level of skepticism towards an m_1 message relative to an m_0 one. Our data shows that this level of skepticism is positively and strongly correlated with buyers' beliefs about the informational content of messages.[48]

5.4 Putting Things Together

We finish this section by looking at the informativeness of messages in the presence and absence of competition. As defined in Section 2.4, informativeness of messages is the difference in beliefs upon observing an m_1 and an m_0 message, and it tells us how much more likely message m_1 is to be sent by a seller with a high-quality product quality than message m_0. Figure 6 presents the *perceived*

[48] In particular, for each buyer in the Competition treatment, we measure the average likelihood of choosing a seller with an m_1 message over an m_0 one separately for the first 5 and the last 5 blocks of the experiment. We then compute the correlation between this measure and the difference in beliefs associated with an m_1 and an m_0 messages, separately for the first 5 and the last 5 blocks. The correlation is 0.57 ($p = 0.004$) for the first 5 blocks and 0.41 ($p = 0.053$) for the last 5 blocks.

informativeness, which captures buyers' beliefs, and the *actual informativeness*, which depicts the correct market interpretation of messages given sellers' behavior.

Figure 6: Perceived and Actual Informativeness of Messages, dynamics.

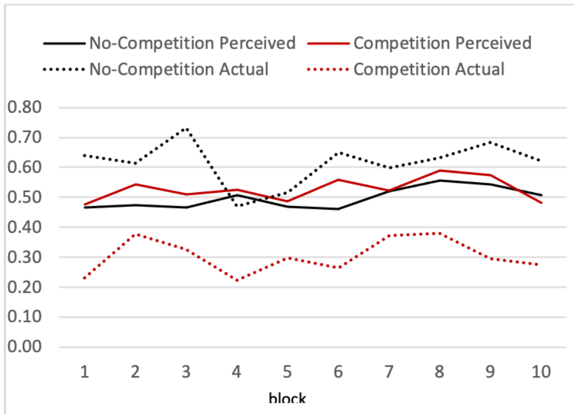

Notes: The solid lines depict the difference between Buyers' beliefs about message m_1 and message m_0 averaged in a block in each treatment. The dotted lines are the difference between probabilities that the product is high quality conditional on message m_1 versus m_0 for actual realized trades.

The comparison between treatments is quite stark and corroborates the mechanism identified above. In the game without competition, buyers' beliefs are close to the actual messages' meaning given sellers' strategies and realized trades in the market. Although a statistically significant difference exists between perceived and actual informativeness in a few initial blocks of the experiment, this difference by and large disappears in the later blocks.[49] By contrast, in the game with competition, the gap between what buyers think messages mean and what they actually mean is large. This gap is persistent and is not mitigated by learning.[50] In other words, competition among sellers diminishes the informational content of messages, but buyers do not realize this is happening and continue to trust messages more than they should.

Bringing all pieces of the story together, we find that markets without competition feature correct average buyer beliefs from the start and higher buyer responsiveness to sellers' strategies via the feedback they observe. In these markets, beliefs are the main driving force of buyers' purchasing decisions, and even buyers who hold initially wrong beliefs adjust them by responding to what transpires in the market. On the contrary, buyers in the markets with competition mistakenly believe from the start of the experiment that the m_1 messages are associated with a high average product quality and the follow-up experience does not convince them otherwise. In general, both

[49]We compare the distribution of perceived informativeness in the No Competition treatment, estimated for each buyer in a block, with the average actual informativeness in the same block and obtain the following $p-$values for each block: $p = 0.014$, $p = 0.005$, $p < 0.001$, $p = 0.517$, $p = 0.423$, $p = 0.008$, $p = 0.229$, $p = 0.180$, $p = 0.038$, $p = 0.091$.

[50]We perform the same statistical analysis for the Competition treatment as the one reported in footnote 49 and obtain the following $p-$values for each block: $p < 0.001$, $p = 0.073$, $p = 0.150$, $p = 0.004$, $p = 0.107$, $p < 0.001$, $p = 0.025$, $p < 0.001$, $p = 0.002$, $p = 0.014$.

beliefs and purchasing cutoffs are more rigid and less responsive to experience in markets with competition, which is why buyers never learn the true meaning of an m_1 message when multiple sellers compete with each other.

We finish by noting that the feedback we provide in our experiment is comparable to what real markets have to offer. In this sense, our results about the sluggishness of buyers' beliefs in markets with competition and the competitive pressure among sellers to lie when they face competition are likely to be present in markets outside the laboratory setting. Indeed, the market structure affects how much buyers and sellers can learn from the feedback they observe. In both markets with and without competition, the sellers can learn the buyers' behavior pretty well. This is consistent with our observation that the sellers' beliefs about buyers' interpretation of messages are on average correct in both markets. Moreover, in markets without competition, the buyers learn everything there is to learn about the sellers except for their psychological types, which allows them to calibrate their beliefs and anticipate the average quality of products with different labels. The situation is different in markets with competition because buyers only learn things about the selected seller and have a limited understanding of the strategy used by the seller they chose not to interact with. Our data shows that this results in buyers' inability to correctly anticipate what messages mean and persistence in buyers believing that competition between sellers is beneficial for them.

6 Conclusions

In this paper, we study the impact of introducing both psychological payoffs and competition into a communication (market) game and investigate their consequences for market outcomes and welfare. Specifically, we look at sellers who suffer a cost when they lie and/or mislead buyers into purchasing subpar goods, and buyers who suffer from disappointment whenever they are tricked into buying such goods.

In contrast to previous experimental work on psychological games, in an effort to construct an equilibrium model and test it in the lab, we induce the costs of lying, guilt, and disappointment, in addition to the material payoffs of our subjects. Doing so allows us to control these psychological payoffs experimentally to some extent and evaluate their comparative static effects.[51]

We find that the introduction of psychological payoffs increases trade and marginally increases the quality of purchased goods. The introduction of competition between sellers, however, undoes these benefits and leads to lower welfare for both buyers and sellers. While in such markets more goods are sold, considerably more of these goods are of low quality.

We analyze the mechanisms underlying these aggregate results and identify the main forces preventing competition from curbing the lying in these markets. We find that competitive pressures, especially in a winner-take-all situation such as ours, encourage sellers to misrepresent and lie more, even if they suffer from the psychological costs of doing so. The sellers' propensity to lie more is reinforced by the behavior of our buyers who fail to understand the lies the sellers tell them in such a competitive environment. In fact, our buyers seem to genuinely believe competition is beneficial

[51]As we have shown, our experimental design was successful in manipulating subjects' payoffs even if one assumes they arrive in the lab with their own aversion to lying and guilt. This is evident from the fact that in our treatments with induced psychological costs, we observe what appear to be partially informative equilibria that exist only when sellers experience psychological costs. At the same time, in the treatment in which we eliminate psychological costs and induce only the material ones, we do not see partially informative equilibria being played. In other words, the presence of psychological costs is beneficial in markets without competition, because they facilitate trade, which does not occur in the pooling equilibria of markets with only monetary payoffs.

for their welfare because they fail to change the way they interpret messages when competition is introduced. The abundant feedback and experience we provide our subjects do not correct for the misperception of messages our buyers demonstrate in the game with competition. Consequently, sellers take advantage of such blind faith on the buyers' part and peddle lower-quality products indiscriminately.

One natural question that arises is what type of intervention may remedy the deleterious impact of competition in these markets. An obvious one is sellers' reputations which were not a feature of our experiment but do exist in the real world. Such reputations are developed in settings where interactions are repeated and where buyers can observe the outcomes of other buyers and their experiences. This should help buyers learn the product quality of sellers they have not bought from and hence add valuable information. Such information may be available on the web via customer reviews which may be of great assistance in steering these markets in the right direction and policing the behavior of those sellers who are less burdened by lying and guilt aversion.[52]

References

Abeler, J., D. Nosenzo, and C. Raymond (2019): "Preferences for Truth-Telling," *Econometrica*, 87(4), 1115–1153.

Battigalli, P. and M. Dufwenberg (2007): "Guilt in Games," *American Economic Review*, 97(2), 170–176.

——— (2009): "Dynamic psychological games," *Journal of Economic Theory*, 144, 1–35.

——— (2022): "Belief-Dependent Motivations and Psychological Game Theory," *Journal of Economic Literature*, 60(3), 833–882.

Born, A. (2020): "Promise Competition," .

Cabral, l. and A. Hortacsu (2010): "The Dynamics of Seller Reputation: Evidence from eBay," *Journal of Industrial Economics*, 58(1), 54–78.

Casella, A., N. Kartik, L. Sanchez, and S. Turban (2018): "Communication in Context: Interpreting Promises in an Experiment on Competition and Trust," *Proceedings of the National Academy of Sciences*, 115(5), 933–938.

Charness, G. and M. Dufwenberg (2006): "Promises and Partnership," *Econometrica*, 74, 1579–1601.

——— (2011): "Participation," *American Economic Review*, 101, 1213–1239.

Chen, Y., N. Kartik, and J. Sobel (2008): "Selecting Cheap-Talk Equilibria," *Econometrica*, 76(1), 117–36.

Costa-Gomez, M. and G. Weizsacker (2008): "Stated Beliefs and Play in Normal-Form Games," *Review of Economic Studies*, 75(3), 729–762.

Crawford, V. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50(6)), 1431–1451.

Dal Bo, P. and G. Frechette (2019): "Strategy Choice in the Infinitely Repeated Prisoner's Dilemma," *American Economic Review*, 109(11), 3929–3952.

Danz, D., L. Vesterlund, and A. Wilson (2021): "Belief Elicitation and Behavioral Incentive Compatibility," *American Economic Review*, 112(9), 2851–2883.

Erat, S. and U. Gneezy (2012): "White Lies," *Management Science*, 58(4), 723–733.

Fischbacher, U. (2007): "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10(2), 171–178.

[52]Cabral and Hortacsu (2010) found that eBay sellers' reputations are instrumental in weeding out bad actors.

Fischbacher, U., C. Fong, and E. Fehr (2009): "Fairness, errors and the power of competition," *Journal of Economic Behavior Organization*, 72, 527–545.

Geanakoplos, J., D. Pearce, and E. Stacchetti (1989): "Psychological games and sequential rationality," *Games and Economic Behavior*, 1, 60–79.

Gillen, B., E. Snowberg, and L. Yariv (2019): "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study," *Journal of Political Economy*, 127(4), 1826–1863.

Gneezy, U. (2005): "Deception: The Role of Consequences," *American Economic Review*, 95(1), 384–394.

Gneezy, U. and J. Potters (1997): "An Experiment on Risk Taking and Evaluation Periods," *The Quarterly Journal of Economics*, 112(2), 631–645.

Goeree, J. and J. Zhang (2014): "Communication Competition," *Experimental Economics*, 17(3), 421–438.

Huck, S., G. Lunser, and J.-R. Tyran (2012): "Competition fosters trust," *Games and Economic Behavior*, 76, 195–209.

Hurkens, S. and N. Kartik (2009): "Would I lie to you? On social preferences and lying aversion," *Experimental Economics*, 12, 180–192.

Jin, G., M. Luca, and D. Martin (2021): "Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure," *American Economic Journal: Microeconomics*, 13(2), 141–173.

Kartik, N. (2009): "Strategic Communication with Lying Costs," *The Review of Economic Studies*, 76(4), 1359–1395.

Keck, S. and N. Karelaia (2012): "Does competition foster trust?" *Experimental Economics*, 15, 204–228.

Nyarko, Y. and A. Schotter (2002): "An Experimental Study of Belief Learning using Elicited Beliefs," *Econometrica*, 70(3), 971–1005.

Sanchez-Pages, S. and M. Vorsatz (2007): "Enjoy The Silence: An Experiment on Truth-Telling," *Experimental Economics*, 12(2).

Schotter, A., A. Weiss, and I. Zapater (1996): "Fairness and survival in ultimatum and dictatorship games," *Journal of Economic Behavior and Organization*, 31(1), 37–56.

Smith, V. (1976): "Experimental Economics: Induced Value Theory," *American Economic Review*, 66(2), 274–279.

Sobel, J. (2020): "Lying and Deception in Games," *Journal of Political Economy*, 128(3), 907–947.

Vanberg, C. (2008): "Why Do People Keep Their Promises? an Experimental Test of Two Explanations," *Econometrica*, 76(6), 1467–1480.

Vespa, E. and A. Wilson (2016): "Communication with multiple senders: An experiment," *Quantitative Economics*, 7, 1–36.