Trust Me: Communication and Competition in a Psychological Game*

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

June 23, 2022

Abstract

We study theoretically and experimentally a communication game with and without sender competition and embed it in a psychological-game framework where senders and receivers incur psychological costs for lying and deceiving. We derive the equilibrium predictions of this model and test them in the controlled laboratory experiment. We find that while the introduction of psychological costs is welfare increasing the introduction of competition is not. The latter result is driven by inordinate amounts 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 have an informed principal who sends a message to a less informed agent 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 the principal to the agent as a function of the divergence of their preferences over material outcomes. However, in many situations, these preferences are known (e.g. that the mechanic can make more money by ripping off her customers). What is not known is how honest the principal is or his degree of guilt aversion. [1]

In such a world, after receiving a message, an agent should consider not only the principal's preferences over material outcomes but also his moral character. If the agent knows the principal is extremely averse to lying or guilt, trusting his messages can become considerably easier. Unfortunately, the agent rarely knows the moral character of a random principal, which is why she must infer it from received messages. Furthermore, such an interpretation is likely to depend on whether

*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] For instance, principals may feel guilty when they knowingly send a misleading message that alters the agent's action and disappoints him, or they may simply have an aversion to deceiving an unsuspecting agent on principle.

the Principal acts in isolation or, on the contrary, competes with other principals for the agent's attention.

Motivated by the above discussion, in this paper, we study the principal-agent game in which subjects have psychological payoffs that depend on messages and beliefs about their reliability, in addition to the standard monetary payoffs. We then introduce competition between two principals who both send messages in an effort to win over the agent. We study the welfare consequences of competition and document whether it acts to curb immoral behavior. We approach this setting from two perspectives: first, by studying a theoretical model of such a psychological-communication game with and without competition and solving for its equilibria, and second, by bringing it to the lab and testing its predictions in a controlled environment.

Setting. We study two games. The first, a game without competition, involves two players: the principal (the seller) and the agent (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 the 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 addition to the material payoffs, players receive psychological payoffs, which depend on their types. Although psychological payoffs can be multi-dimensional, we focus on guilt and lying aversion for the sellers and disappointment for the buyers. Sellers can suffer from lying when they misrepresent the quality of the good they own, and can also feel guilty if they mislead the buyers about the quality of their good. Buyers can suffer from disappointment when their expectations about the quality of the good are misplaced. In the second game—a game with competition—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.

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 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 the same or a more informative equilibrium is played in the presence of competition, buyers stand to benefit from it because they are more likely to select a seller with a high-quality product if messages are as informative or more informative in the game with competition. However, if sellers with low-quality products feel compelled to lie more due to competition, they 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 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.

Experimental Results. Whereas expected welfare increases when psychological payoffs are introduced, we find that welfare for both sellers and buyers is unambiguously lower in the game with competition. Three reasons explain this finding. First, subjects do not play the pooling equilibrium in either of the two games. Instead, they play as if seller's messages transmit information that buyers can rely on.[2] Second, sellers with low-quality products tend to lie more when they compete with each other; this pattern is observed even for those sellers who suffer from high lying costs and are highly guilt averse.[3] Third, buyers fail to adequately adjust their beliefs to the increased dissembling of sellers, indicating they do not best respond to the strategies that sellers play in the game with competition. In fact, buyer beliefs and purchasing decisions are remarkably similar in the two games. Therefore, from a buyer's perspective, the game with competition features an additional benefit of more sellers to choose from, which leads them to purchase goods more often in the game with competition even though they end up with a low-quality product more often. Consequently, the average buyer payoff is lower in the presence of competition. The average seller payoffs are also lower in the presence of competition, because of the higher psychological costs they incur when they own a low-quality product and try to convince the buyer otherwise.

Contribution and connection to the literature. Our theoretical model builds on two branches of literature: psychological games that incorporate belief-dependent preferences (Geanakoplos, Pearce, and Stacchetti (1989), Battigali and Dufwenberg (2007, 2009)) and games with lying costs (Chen et al (2008), Kartik (2009), Sobel (2020)). As we mentioned earlier, principals in our model may experience guilt if they lead the agent 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 agent's actions or beliefs. The agent in our model can experience disappointment when she relies on the principal'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 (Charness and Dufwenberg (2006, 2011), Vanberg (2008), Goeree and Zhang (2014), and Casella et al (2018)). 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.[4] The induced value approach has been successfully implemented in a variety of individual-decision

[2]We also conduct an additional treatment with induced material payoffs only and find that in this treatment, the majority of buyers and sellers adhere to pooling equilibrium strategies. This finding shows the introduction of psychological payoffs is indeed responsible for shifting game play from the uninformative equilibrium to the partially informative one. Moreover, the comparison between this treatment and the one with psychological payoffs and no competition between sellers provides additional validation for our method of inducing psychological payoffs alongside the material payoffs.

[3]The idea here is that whereas, in equilibrium, none of these subjects should lie, the possibility of not being chosen for a sale leads them to lie more often despite their lying and guilt aversion.

[4]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.

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.[5] The three most closely related papers to ours are Casella, Kartik, Sanchez, and Turban (2017), Goeree and Zhang (2014), and Born (2018). Casella et al. (2017) 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. (2017), we find the shift in the communication strategies when competition is present, but contrary to Casella et al. (2017), 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.[6]

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 (2012) 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 (2018) 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 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 a stable behavior.

Structure of the paper. We proceed as follows. In section 2, we introduce the communication game with psychological payoffs, with and without competition between sellers, and solve for its

[5] Several studies look at the effects of competition on trust in various environments (Huck et al. (2012), Keck and Krelaia (2012), and 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.

[6] Similar results are found by Jin, Luca, and Martin (2017) 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 adjust for it.

equilibria. This game serves as the basis for the experiment we run. In section 3, we describe the experimental design and its implementation. Section 4 contains results of the experiment, and section 5 offers some conclusions.

2 The Model

We study a communication game between an informed Seller (he) and an uniformed 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 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 not, and the game ends.

Despite the simplicity of this description, our game gets complicated by the fact that the Seller's and Buyer's payoffs are influenced by psychological payoffs that depend on their feelings of guilt and lying aversion (for sellers) and disappointment (for buyers). Both players know such payoffs exist and must take them into account when evaluating the veracity of messages.

We study this game in two contexts: the no-competition case, which has only one Seller and one Buyer, and the competition case, which has two Sellers competing with each other to sell to a single Buyer. In the latter case, the Buyer chooses one of the Sellers to deal with. 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 Material Payoffs

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.

Absent any psychological payoffs, the game has a unique Bayesian Nash equilibrium (equilibrium hereafter), in which trade does not occur. 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 message m_i, which contradicts our initial presumption. Thus, no equilibrium can exist in which one message entails a higher probability of the high-quality product than another. Therefore, the Buyer is left with her prior beliefs, and given the material payoffs, *the no-trade equilibrium is the only one that exists.*

Figure 1: Material Payoffs in the Game without Competition

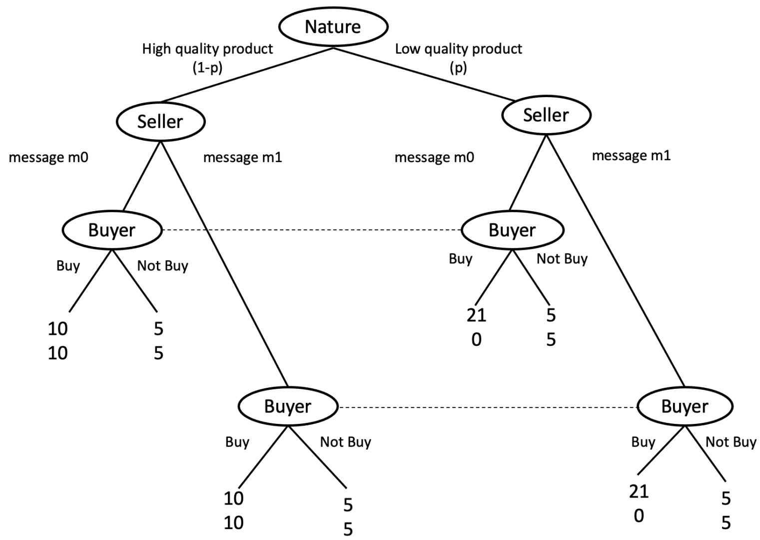

Notes: The top payoff at each node is the Seller's material payoff and the bottom one is the Buyer's material payoff. The dashed line indicates the information set of the Buyer because she does not know the type of Seller she is dealing with.

2.2 Psychological Payoffs

The above analysis changes when we introduce psychological payoffs. Players are now motivated not only by their material payoffs, but also by belief-dependent utilities that are determined by the communication strategies used by the Seller and by the players' first- and second-order beliefs about each others' actions. Following the literature, we focus on several psychological forces that have previously been identified as important.[7] We present here the main ingredients of our behavioral model and refer the reader to the Online Appendix section 1 for the 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

[7] See Dufwenberg and Gneezy (2000), Gneezy (2005), Battigalli and Dufwenberg (2007), Vanberg (2008), Charness and Dufwenberg (2006, 2011), and Ellingsen et al. (2010). In section 2.7, we discuss our choice of psychological payoffs and its connection to the literature.

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.[8] The Seller chooses a decision function, which 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,

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

where $s^B(\omega, m_i)$ stands for the likelihood that the Buyer with type ω purchases the product after observing message $m_i \in M$.

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

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

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.[9] Recall that the two messages are $m_0 =$"The product is low quality" and $m_1 =$"The product is really high quality".

The lying component of the Seller's psychological payoff is determined by the difference between the message the Seller sends and the quality of the good he possesses. In other words, lying aversion captures moral objection to lying. 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. 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

[8]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 Online Appendix section 1).

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

sells her the low-quality product. Formally, the amount of guilt that the Seller experiences is equal to

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

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. We denote the last term as $\mathbb{E}[\omega|\text{Buyer with type } \omega \text{ buys the product after receiving } m_1]$ by $\mathbb{E}[\omega \text{ when Buy}|m_1]$.

Finally, the Buyer may feel disappointed after believing the Seller's claim of a high-quality product and ending up buying the low-quality one. We define this disappointment as equal to the difference between the expected and actual material payoff the Buyer receives conditional on observing m_1; that is,

$$10 \cdot b_B^1(m_1) - 0.$$

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 a 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.[10]

2.3 Equilibrium Analysis of the Game without Competition

Equilibrium consists of specifying a decision function for the Seller, s^S, indicating the probability distribution over messages for each product quality and the Seller's psychological type, a decision function for the Buyer, s^B, indicating the probability that the Buyer buys the product for each message and the Buyer's psychological type, and the system of beliefs of both the Buyer and the Seller (b_B^1, b_S^2) such that (a) the Buyer's purchasing decisions are optimal given the decision function of the Seller and Buyer regarding the Seller's messages, (b) the Seller's messages are optimal in the sense that no Seller wants to send a message different from the one specified by his decision function given his beliefs, and (c) beliefs are "correct"; that is, first- and second-order beliefs of players coincide with the expected frequency of the Buyer choosing to purchase the product conditional on received messages.

In the analysis that follows, we focus on equilibria in which $b_B^1(m_0) = b_S^2(m_0) \leq b_B^1(m_1) = b_S^2(m_1)$. We view this natural restriction as capturing the idea that a message, m_1, that states the

[10] Because the Seller does not know the value of the Buyer's ω, the actual payoff to the seller is

$$21 - 10 \cdot G \cdot b_S^2(m_1) \cdot \mathbb{E}[\omega \text{ when Buy}|m_1] \cdot 1_{m_i=m_1} - L \cdot 1_{m_i=m_1} \quad \text{if } q = q_L,$$

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

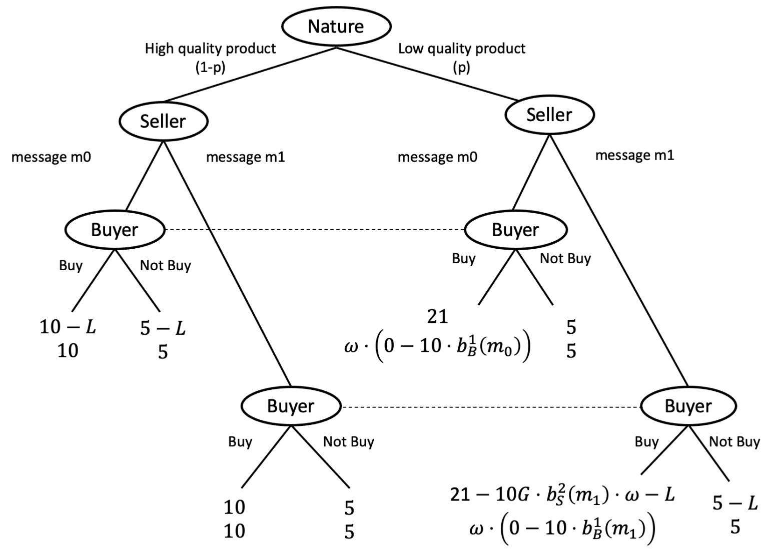

Notes: The top payoff at each node is the Seller's overall payoff, and the bottom one is the Buyer's overall payoff. The dashed line indicates the information set of the Buyer because she does not know the type of Seller she is dealing with.

product is of high quality implies a weakly higher belief regarding the product quality than the other message, m_0.[11] Also, we we often write $s^B(\omega, m_i) =$ Not Buy instead of $s^B(\omega, m_i) = 0$ and $s^B(\omega, m_i) =$ Buy instead of $s^B(\omega, m_i) = 1$ when this creates no confusion.

As standard in any cheap-talk game, a non-informative babbling equilibrium exists in our setting as well, in which messages sent by the Seller convey no information, and therefore, the Buyer chooses the same distribution of actions regardless of the message. We refer to this equilibrium as the pooling equilibrium and note that in our setting, a **unique pooling equilibrium** exists in which the Buyer never buys the product and secures the payoff of 5 (no-trade equilibrium).

This pooling equilibrium with psychological payoffs coincides with the unique Bayesian Nash equilibrium of the game with just material payoffs. The pooling equilibrium, however, is not the only type of equilibrium that one can sustain in our psychological game. Under some restrictions on the game's primitives, **partially informative equilibria (PIE)** exist in which some of the information about the product quality is conveyed in the communication stage between the Seller and the Buyer. These equilibria rely 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, for a given message m_i and belief $b_B^1(m_i)$ associated with it, the Buyer

[11] Absent this restriction, additional equilibria exist in which the meanings of messages are flipped; that is, message m_1 is interpreted as the product being a low quality, and m_0 is interpreted as the product being a high quality.

will choose to buy the product if and only if

$$\mathbb{E}\Pi^{\text{Buyer}}\left(m_i, \text{Buy}\right) \geq \mathbb{E}\Pi^{\text{Buyer}}\left(m_i, \text{Not Buy}\right) \Leftrightarrow$$

$$10 \cdot b_B^1(m_i) + (-10 \cdot b_B^1(m_i) \cdot \omega) \cdot (1 - b_B^1(m_i) \geq 5 \Leftrightarrow$$

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

where $\bar{\omega}(m_i)$ denotes the Buyer's type who is indifferent between purchasing and not purchasing the product after observing message m_i.

Further, Sellers with high-quality products send the message m_1 irrespective of their psychological type, because (a) lying is costly and (b) $b_B^1(m_1) \geq b_B^1(m_0)$, which means that more Buyer types will choose to buy the product because $\frac{\partial \bar{\omega}(m_i)}{\partial b_B^1(m_i)} > 0$. Thus, message m_0 necessarily comes from the Seller with a low-quality product, which implies $b_1^B(m_0) = b_2^S(m_0) = 0$, and the Buyer does not purchase the product after observing message m_0, implying $s^B(\omega, m_0) = \text{Not Buy}$ for any ω.

The behavior of Sellers with low-quality goods depends on her psychological type. Sellers with relatively lower values of guilt and lying sensitivities may lie and announce they have high-quality products, whereas others with higher values of either guilt or lying sensitivities will tell the truth about the product quality.[12] Thus, in any PIE, $b_B^1(m_1) = b_S^2(m_1) > 0$. Further, denote by $H(\cdot)$ the distribution of disappointment-sensitivity parameters of the Buyers. Then, the Seller with a low-quality product and type (G, L) prefers to be truthful and send message m_0 if and only if

$$\mathbb{E}\Pi^{\text{Seller}}\left(m_0, \text{Not Buy}\right) \geq \mathbb{E}\Pi^{\text{Seller}}\left(m_1, s^B(m_1)\right) \Leftrightarrow$$

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

Finally, in equilibrium, beliefs must be correct, namely,

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

where ψ is the proportion of Sellers with a low-quality product 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.

2.4 The Game with Competition

We now introduce competition between Sellers. Two sellers (Seller 1 and Seller 2) compete for the opportunity to sell the good to a single Buyer.

This game 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

[12] If a Seller exists who owns a low-quality product and has psychological type $(0, 0)$, that is, a Seller with a low-quality product who is motivated only by material payoffs, he will necessarily send message m_1 in any PIE.

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.

Our analysis focuses on symmetric equilibria, in which both Sellers use the same decision functions when they are of the same psychological type, own a product of the same quality, and hold the same beliefs. The **symmetric equilibrium** in this game consists of specifying the following: the decision function for the Sellers s^S, indicating the probability distribution over messages for each Seller's psychological type and each product quality; the selection function for the Buyer, indicating which Seller is selected based on observed messages; and a decision function for the Buyer s^B, indicating the probability that the Buyer buys the product for each message received from the selected Seller. The equilibrium also includes the system of beliefs for both the Buyer and two Sellers (b_B^1, b_S^2) such that (a) the actions of the Buyer are optimal given her beliefs and the decision function of the Sellers (both in terms of which Seller she selects to play the tree game with and her actions in the tree game), (b) the decision function of the Sellers is optimal given their beliefs and the actions of the Buyer, and (c) all beliefs are correct.

In general, the game with competition admits the *same types* of equilibria as the game without competition. Specifically, a **unique pooling no-trade equilibrium** always exists in which the Buyer treats all messages received from the Sellers as uninformative and randomly selects one of the two Sellers to buy from. In this equilibrium, the selected Seller and the Buyer get payoffs of 5 each, whereas the unmatched Seller gets a payoff of zero. In addition, one can also support a PIE in our game with competition. In this case, 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 game with competition and the game without competition both have the same sets of equilibria: a unique pooling equilibrium, and two PIEs, which differ in the fraction of low-quality Sellers who lie in the communication stage. The choice of these parameters was deliberate because they allow us to investigate the equilibrium selection process that occurs in the presence and absence of competition and to determine whether the competition among Sellers is beneficial or detrimental for the Buyers.

Although we establish the existence of an equilibrium for our model using the parameters employed in our experiment, in the Online Appendix, we provide general conditions under which, in addition to the always existent uninformative pooling equilibrium, a partially informative equilibrium (PIE) also exists in which the messages sent by our Sellers have some informational content.

2.5 Parameterization

To bring our model to the lab, we need to set parameters to be used in both games. We choose the following values. The probability that the product quality is low is $p = 60\%$, which ensures that, absent psychological utilities, both games have the unique pooling equilibrium where the Buyer never purchases the product. The distribution of 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: $(L, G) \in \{(0, 0), (0, 6), (20, 0), (20, 6)\}$.[13] 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 $\Pi^{\text{Seller}}(m_0, \text{Not Buy}) = 5$. This modification of payoffs does not affect the qualitative predictions of the theory.[14]

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, the Buyer treats messages from the Seller(s) as uninformative and does not update her prior beliefs about the Seller's 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,

[13] 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.

[14] The theoretical predictions presented here incorporate this modification.

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.

2.6 General Theoretical Insights

The three equilibria described above are ranked in terms of how much information Sellers transmit in the communication stage. The pooling equilibrium is the one with the least amount of information because messages are not informative and the Buyer's posterior beliefs about the quality of the Seller after observing either of the two messages is identical to the prior. In PIE1 and PIE2, messages are partially informative because they change the posterior probabilities that Buyers assign to the Sellers' types. In our setup, PIE2 is the most informative equilibrium one can support, because the only psychological type of the low-quality Seller who lies in this equilibrium is the type who suffers no psychological disutility from lying and guilt. We define the notion of the informativeness of an equilibrium as the difference between the posterior belief of the Buyer after observing message m_1 and message m_0 and denote it by Eq^{info}. That is,

$$Eq^{\text{info}} = b_B^1(m_1) - b_B^1(m_0).$$

The larger this difference, the more information the Buyer learns from the Sellers' messages. The next observation summarizes the discussion regarding informativeness of different types of equilibria in both versions of the game.

Observation 1. *The three equilibria discussed above are ranked in terms of their informativeness, with pooling equilibrium being the least informative and PIE2 being the most informative*

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

Informativeness of equilibria can directly be translated into the Buyer's expected payoffs. The more information the Buyer receives from the Sellers' messages, the better purchasing decisions she can make. This is summarized in the next observation, which again applies to both versions of the game.

Observation 2. *In both games with and without competition, the expected payoff of the Buyer monotonically increases with the informativeness of the equilibria; that is,*

$$\mathbb{E}\Pi^{Buyer}_{POOL} < \mathbb{E}\Pi^{Buyer}_{PIE1} < \mathbb{E}\Pi^{Buyer}_{PIE2}.$$

Given the multiplicity of equilibria in both versions of the game, understanding whether competition among the Sellers is beneficial or detrimental to the Buyer is a tricky question. In particular, the answer would depend on which equilibrium is selected in each version of the game, because different equilibria result in different expected Buyers' payoffs. This question is ultimately an empirical one, and a compelling reason to conduct laboratory experiments. Controlled and carefully designed experiments can inform us about equilibria selection issues, which are extremely hard to disentangle using data collected outside the lab.

When the same equilibrium is played in the two games, however, one can compare Buyers' expected payoffs, as we show in the observation below.

Observation 3: *If the same equilibrium is played in games with and without competition, the Buyer's expected payoff is weakly higher in the game with competition.*

The logic behind Observation 3 is straightforward: when messages are not informative, as in the case of the pooling equilibrium, no scope exists for competition to influence Buyers' payoffs. The benefit of competition between Sellers comes into play only when messages are partially informative as in PIE. If the same PIE is played in both games, Sellers' decision functions that determine messages that different types send are the same. In this case, the presence of two Sellers increases the likelihood that one of the Sellers receives a high-quality good, sends message m_1, and is selected by the Buyer. This is clearly beneficial to the Buyer because it increases the chances of the Buyer playing the game with the high-quality Seller.

However, the Buyers' payoffs might decrease with competition for several reasons. First, a Seller might be compelled to lie more often when facing competition when he has a low-quality product. Indeed, if he reveals the quality truthfully, the chance that the other Seller will make the sale instead is greater, either because the latter has a high-quality good and sends the message m_1, or because the latter has a low-quality good and still sends message m_1 (lies!).[15] Thus, in the face of competition, we may see more lying and hence more opportunities for Buyers to purchase low-quality goods. Second, the introduction of competition could lead to the selection of different equilibria, some of which have lower Buyers' payoffs. Given these reasons, an empirical investigation is required to assert whether the competition among Sellers is ultimately beneficial or detrimental for the Buyers.

2.7 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 Sellers' products nor the type of 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

[15] Schotter, Weiss, and Zapater (1996) have shown competition can have an impact on rejection behavior in ultimatum games 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.

paper by Goeree and Zhang (2014), who introduce competition between Sellers and find communication and competition act as substitutes. Although our game shares some similarities with the game in the above-mentioned paper, important differences also exist.[16] 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 (5 in our game). This adds an element of reciprocity on the part of 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 the simplest communication game with hidden information, which is amenable to psychological payoffs and the introduction of competition among Sellers.

Psychological payoffs. Our behavioral model incorporates several psychological motives, including guilt, lying aversion, and disappointment. Although this list is obviously not exhaustive, the motives we study here are among the natural candidates identified by the literature in various communication games (see Charness and Dufwenberg (2006, 2001), Vanberg (2008), Goeree and Zhang (2014), Casella et. al. (2018)). In the last decade, both theoretical and experimental literature have made significant progress in carefully defining various psychological costs and differentiating between them. We contribute to this literature by studying equilibrium effects of introducing such costs in a communication game with hidden information and its interaction with competition. Therefore, describing the connection between our psychological forces and those used in the literature is useful. We offer this description next.

Our definition of ***lying*** is consistent with the notion used in Kartik (2009) and Sobel (2020). Lying means saying things that are not true, that is, misrepresenting your private information. 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.[17] In our model, Sellers dislike lying and suffer a cost when they claim their low-quality product is of high quality. This approach is consistent with the large experimental literature that documents that people lie surprisingly little relative to what theory predicts they should if they only cared about material payoffs. Abeler et. al. (2019) make this point by providing a meta-data analysis of experimental work on lying and showing that lying aversion is driven by a preference to be seen as honest and a preference to be honest per se.

The two other psychological costs, guilt and disappointment, depend on players' beliefs. For the Seller, the amount of ***guilt*** he experiences depends on the Buyer's interpretation of messages and on how the Buyer acts on it. This concept is similar to the notion of deception in Sobel (2020),

[16] 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 that second game than in the first one.

[17] A Seller can send one of the two messages: $m_0 =$"The product is low quality" and $m_1 =$"The product is really high quality."

where deception depends on how the receiver interprets messages and how her actions might change in response to them.[18] Our definition of guilt also relates to Battigali and Dufwenberg's (2007) concept of guilt in games, which captures a failure to live up to receivers' expectations. 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. Moreover, we allow Buyers to experience ***disappointment*** if their expectations about product quality do not square with the product they actually receive conditional on a message. We believe the introduction of Buyers' disappointment is natural in the game we study. 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.[19]

Finally, we note that in the equilibrium of our model, both guilt and lying aversion prevent some low-quality Sellers from sending a misleading message to a Buyer. However, the two forces play out differently. To illustrate, we re-examine the equation that determines whether 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:

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

, where $\bar{\omega}(m_1)$ is the Buyer type that is indifferent between buying the product and not after observing message m_1. The first term on the left captures that low-quality Sellers with high aversion to lying are likely to be truthful in the communication stage. But even absent lying aversion, that is, when $L = 0$, the guilt force alone might prevent some low-quality Sellers from lying, which is the second term on the left. This guilt force takes into account not only the Seller's second-order beliefs about the Buyer's interpretation of message m_1, but also the Buyer's tendency to purchase the product given message m_1, which is incorporated in $\bar{\omega}(m_1)$. The presence of both forces, guilt and lying aversion, makes our game interesting in the sense that in both the game with and without competition, *two* **PIE**s exist with different equilibrium interpretations of messages. Thus, our setting allows us to explore the inflation of language, which might arise in the game with competition and will be detectable by differential responses to messages in the two games.

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

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

[19] 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 parameter. As a consequence, the Seller's guilt is tied to the extent to which he anticipates his actions will disappoint a Buyer he is facing.

3.2, we discuss benefits and challenges of inducing psychological costs in the lab experiment. In section 3.3, we describe how we elicited both actions and subjects' beliefs in an experiment with induced psychological costs.

3.1 The Design

Our experiment is a direct implementation of the model described above. We conducted three separate treatments: a no-competition treatment, a competition treatment, and a monetary treatment. Whereas the first two treatments incorporate psychological payoffs, the last one does not and essentially is the situation depicted in Figure 1. One unique innovation of our experiment is that we induce the psychological payoffs described in Figure 2 for the no-competition treatment and the competition treatment. So we impose costs on the Seller whenever he lies to the Buyer and disappoints her. We also impose a sensitivity parameter ω on the Buyer that specifies her sensitivity about being misled by the Seller. As discussed above, these lying and guilt costs (L and G) are induced and take on different values depending on the type of the Seller, 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.[20]

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 the Supplementary 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. Each Seller wants to sell his product to the Buyer. The chance that the product is of high quality is 40%, and the chance that it is of low quality is 60%. The Seller always knows the quality of his product but the Buyer does not. Each Seller can send a message $\{m_{0,}m_1\}$ to the Buyer to convince her to buy the product. The Buyer has to decide whether to buy it, based on the message she receives.

The Seller can be one of the four types with equal probability: Type S1 has $(L, G) = (0, 0)$, Type S2 has $(L, G) = (0, 6)$, Type S3 has $(L, G) = (20, 0)$, and Type S4 has $(L, G) = (20, 6)$. The disappointment parameter for the Buyer, ω, is drawn randomly and uniformly over the interval $[0, 1]$. 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 that asks them to specify the messages they want to send conditional on the quality of their goods and their randomly determined types (see Figure 3).

The task of the Buyer 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

[20] Given that we induce psychological payoffs, our focus is not on assessing whether real-world agents suffer from guilt 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

is less than $\omega'(m_0)$ ($\omega'(m_1)$), the Buyer buys the good (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 were also asked to enter their beliefs. Each Buyer was asked to enter a number between 0 and 100 representing her belief that the Sellers who sent 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 block, subjects maintained their roles but were randomly assigned new partners.[21]

We use this block design because entering a strategy and a set of beliefs and reviewing feedback is a time-consuming process, and hence would be very time consuming for subjects 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

[21]The screenshots depicting feedback that subjects received at the end of each block are presented in the Supplementary Appendix.

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.[22] 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 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 **??** 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 $R = 2.5$, 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, Snowberg, and Yariv (2018).

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, Sellers were asked to specify the message that would be sent to the Buyer for each possible product quality they might possess, and Buyers were asked to specify their purchasing decision for each of the two messages they could receive from the Seller. We also elicited Buyers' and Sellers' beliefs as before. Our interest in introducing this treatment was to compare behavior between this treatment and the **no-competition** treatment and assess the effect of introducing (and controlling) psychological forces on the market outcomes and strategies of the market participants. Our experimental design is summarized in Table 1.

3.2 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

[22]Dal Bo and Frechette (2016) use a similar method when they study infinitely repeated prisoners' dilemma games and use the strategy method.

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

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 and lying aversion might be 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. Two observations are relevant to this discussion. First, note that from a theoretical point of view, either the penalty we impose for lying 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 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 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 differ 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.

Second, recall that one of the basic theoretical insights of our model is that the introduction of psychological costs might lead to more truthful communication and to higher welfare of the Buyers. The reason is that the game with psychological payoffs admits partially informative equilibria, which cannot be sustained in the game with only material payoffs. Thus, if we observe that the Sellers lie less in the no-competition treatment than in the monetary treatment, this finding would support the insight described above, whether we managed to fully control psychological forces or did so only partially.

In the Results section, we start by comparing the monetary treatment with the no-competition treatment to address this point. We show that without induced psychological costs, Sellers lie to a far greater extent than they do when such costs are induced, 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 then 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 in the market.

3.3 Eliciting beliefs and actions in psychological games

In psychological games, payoffs are a function of both actions and beliefs, which is why isolating the pure incentive for truthful belief revelation is difficult. For example, a Seller can inflate his game payoff by reporting a belief that minimizes the psychological costs included in his game payoffs, by appropriately reporting second-order beliefs about buyers to minimize what they claim to be the expected amount of disappointment they are creating by lying. In other words, by strategically reporting a belief of 0 that the Buyers will believe an m_1 message, a Seller can eliminate his guilt costs regardless of what level his G turns out to be.

To deal with this issue, we have chosen the parameters of the quadratic scoring rule for beliefs' elicitation in such a way that the distance between the true belief and the optimal reported belief is minimal, and misreporting one's true beliefs increases one's payoffs by a small amount. (We also informed subjects that they would be paid either for the performance in the game or for the accuracy of their beliefs to make hedging between the two harder.)

We illustrate this approach for Buyers' beliefs here.[23] Consider a Buyer who reports her belief about the probability that a Seller has a high-quality product conditional on sending message m_i. Denote by p_{m_i} the Buyer's true belief, and by r_{m_i}, the reported belief. 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

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

The Buyer's payoff from playing the game is

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

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.[24]

[23] We refer the reader to the Online Appendix, where we describe this procedure in detail.

[24] In addition, in the instructions, subjects were strongly urged to report beliefs truthfully to maximize their payoff

4 Results

This section is organized as follows. First, we study the impact of the introduction of psychological payoffs in the game without competition. To do so, we compare market outcomes and behavior in the monetary treatment, in which only monetary payoffs are induced, and the no-competition treatment, in which both monetary and psychological payoffs are induced. Second, we study the effect of competition on market outcomes in communication games with psychological payoffs and compare the no-competition with the competition treatment, both of which feature psychological payoffs in addition to the material ones. Third, we identify the key elements driving these aggregate results using our theory as a benchmark against which we compare strategies and beliefs of market participants. In this analysis, we pay special attention to how Sellers use messages and how Buyers interpret them in the presence and absence of competition.

Approach to Data Analysis. As described in section 3.1, the Buyers and the Sellers state their full strategies and beliefs at the beginning of each block for 10 blocks in each treatment. However, these observations are not independent as subjects were re-matched into different groups in each block and interacted with each other. Therefore, throughout this section, we use the 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), where we regress the variable of interest (e.g., purchasing decision of Buyers, or the quality of the sold product, or Buyers' beliefs) 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 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 due to learning and experience in the game.

4.1 The Impact of Psychological Payoffs in the Game without Competition

We start by exploring the impact of psychological payoffs on market outcomes and participant behavior in the game without competition. We address two intertwined questions in this section. The first question is whether our method of inducing psychological payoffs worked. As we discussed in section 3.2, either the psychological costs we impose are greater or smaller than the subject's own self-imposed homegrown penalties. If the former, we are in full control of subjects' payoffs, both psychological and material. If the latter, a subject may experience greater discomfort from lying or guilt than the penalties we impose in our experiment. In that case, we should observe that

in the experiment, which is a standard procedure in lab experiments. See Danz, Vesterlund, and Wilson (2022) for a discussion of beliefs' elicitation methods commonly used in the lab experiments.

Sellers lie weakly less in the monetary treatment than in the no-competition treatment, which differ only in that psychological costs are induced in the latter treatment. Similar frequencies of lying across these two treatments would indicate subjects' self-imposed psychological penalties dominate the ones we imposed, whereas less lying would indicate partial crowding out of induced and innate costs. The second question is whether the introduction of psychological payoffs affects equilibrium selection, that is, whether it is able to shift outcomes from an uninformative pooling equilibrium, which is the unique prediction for the game with only material payoffs, to a more informative one, which can be supported when psychological payoffs are present.

To address these questions, we compare behavior and outcomes in the monetary treatment with the no-competition treatment, starting with the aggregate perspective and then moving on to the individual-level analysis of strategies.

Aggregate Results. Table 2 presents average behavior of our Sellers and Buyers, separately for each treatment and each half of the experiment.

Table 2: Aggregate Behavior of Buyers and Sellers in the Monetary and No-Competition treatments

		Monetary	No Competition	Monetary vs No Competition
first 5 blocks				
Sellers' behavior	$\Pr[m_1 \mid q = q_H]$	0.96 (0.02)	0.90 (0.06)	$p = 0.303$
	$\Pr[m_1 \mid q = q_L]$	0.60 (0.07)	0.28 (0.03)	$p < 0.001$
Buyers' purchasing behavior	$\Pr[\text{Buy} \mid m_0]$	0.08 (0.02)	0.32 (0.04)	$p < 0.001$
	$\Pr[\text{Buy} \mid m_1]$	0.56 (0.07)	0.59 (0.03)	$p = 0.689$
Buyers' beliefs	$\bar{b}_B(m_0)$	0.26 (0.07)	0.24 (0.03)	$p = 0.726$
	$\bar{b}_B(m_1)$	0.60 (0.03)	0.71 (0.01)	$p < 0.001$
last 5 blocks				
Sellers' behavior	$\Pr[m_1 \mid q = q_H]$	0.92 (0.01)	0.89 (0.06)	$p = 0.657$
	$\Pr[m_1 \mid q = q_L]$	0.62 (0.02)	0.24 (0.03)	$p < 0.001$
Buyers' purchasing behavior	$\Pr[\text{Buy} \mid m_0]$	0.13 (0.06)	0.32 (0.03)	$p = 0.012$
	$\Pr[\text{Buy} \mid m_1]$	0.39 (0.04)	0.56 (0.05)	$p = 0.005$
Buyers' beliefs	$\bar{b}_B(m_0)$	0.19 (0.04)	0.27 (0.05)	$p = 0.208$
	$\bar{b}_B(m_1)$	0.55 (0.05)	0.76 (0.02)	$p < 0.001$

Notes: The average quantities are presented with the robust standard errors in parentheses. $\bar{b}^1_B(m_i)$ is the Buyer's first-order belief about the chance that the product is of high quality after observing message m_i.

Starting with behavior of Sellers, we note that Sellers who own the high-quality product almost always send the truthful message, whereas those who own the low-quality product often lie to the Buyer and send message m_1 instead. Importantly, low-quality Sellers lie significantly more often in the monetary treatment than in the no-competition treatment in both halves of the experiment. This effect is the opposite of what would have happened if subjects' self-imposed psychological costs had been more severe than those we imposed as part of our design. We take this evidence as validation of our experimental method of inducing psychological payoffs and move on to investigating whether a pooling equilibrium is a good description of subject behavior and the final outcomes observed in the game without competition.

Theoretically, in the pooling equilibrium, Buyers treat messages as uninformative and therefore hold the same beliefs and act the same way irrespective of the message they receive. In addition, Sellers send messages in such a way that the quality of the good they hold and the message they send are uncorrelated. The data presented in Table 2 show that although the exact predictions of the pooling equilibrium do not hold in either of the treatments, Seller and Buyer behavior are much closer to the pooling equilibrium in the monetary treatment than in the no-competition treatment. Indeed, focusing on the last five blocks in the experiment, we observe that in the monetary treatment, Sellers who own a low-quality product lie more often about it, Buyers hold lower beliefs about product quality when they receive the m_1 message, and as a result, they are less likely to purchase the product when they receive the m_1 message.

These trends are confirmed by the correlation between product quality and messages sent by the Sellers. Although these correlations are significantly different from zero in both treatments, they are *twice as large* in the no-competition treatment as they are in the monetary one. In the last five blocks of the experiment, these correlations are $corr = 0.65$ and $corr = 0.32$ in the no-competition and the monetary treatments, respectively. Put differently, messages are far more informative in the no-competition treatment with psychological payoffs.

From a slightly different perspective, one can ask, given the behavior of Sellers in the two treatments, whether the best response for Buyers is to purchase the product upon observing m_1 in the no-competition treatment and not purchase it in the monetary treatment. This would be another way to explore whether the Sellers were significantly more honest in the no-competition treatment than in the monetary treatment and whether this difference is high enough to warrant Buyers to purchase the product after receiving an m_1 message in the no-competition treatment but not in the monetary treatment. To answer this question, we first calculate the observed probability of the good being of high quality given message m_1 and then use that probability to calculate the expected payoff for subjects. Remember that if a Buyer does not buy, she can guarantee herself a payoff of five, so buying conditional on receiving an m_1 message would have to yield more than 5 in order to be profitable. In the monetary treatment, we obtain $\mathbb{E}\Pi^{\text{Buyer}}_{\text{first 5 blocks}}|m_1 = 5.26 > 5$ and $\mathbb{E}\Pi^{\text{Buyer}}_{\text{last 5 blocks}}|m_1 = 4.93 < 5$, whereas for the no-competition treatment, we get $\mathbb{E}\Pi^{\text{Buyer}}_{\text{first 5 blocks}}|m_1 = 6.58 > 5$ and $\mathbb{E}\Pi^{\text{Buyer}}_{\text{last 5 blocks}}|m_1 = 7.06 > 5$. These calculations show that buying is marginally profitable in the first five blocks but not in the last five blocks in the monetary treatment, whereas it is always profitable in the no-competition treatment. In other words, whereas Buyers learn to be more skeptical of Sellers' messages in the monetary treatment, the opposite happens in the no-competition treatment, in which the m_1 message becomes, if anything, more reliable over time.

Table 3: Trade Frequencies in Monetary and No-Competition treatments

Treatment	first 5 blocks	last 5 blocks	first 5 vs last 5
Monetary	0.45 (0.06)	0.32 (0.04)	$p < 0.001$
No Competition	0.46 (0.02)	0.43 (0.02)	$p = 0.286$
Monetary vs No Competition	$p = 0.834$	$p = 0.009$	

Notes: Average trade frequencies are reported with robust standard errors in parentheses.

Finally, we turn to the overall trade frequencies in each treatment. Table 3 depicts these and shows that although the trade frequencies are similar in the first half of the experiment, with experience, subjects learn to trade less when psychological payoffs are not induced. In the second half of the experiment, Buyers are significantly less likely to purchase the product in the monetary treatment than in the no-competition treatment.

Buyers' and Sellers' Individual Strategies. To analyze individual strategies of Sellers and Buyers, we interpret the pooling equilibrium as follows. For the Sellers, playing a pooling equilibrium means choosing the same message irrespective of the product quality in the monetary treatment and irrespective of both the product quality and one's psychological type in the no-competition treatment. Adhering to a pooling equilibrium for the Buyers entails making the same purchasing decision no matter what message is received in the monetary treatment, and choosing the two "similar" purchasing cutoffs conditional on either message in the no-competition treatment, where similar means cutoffs are no more than 0.05 away from each other.[25] Table 4 presents the frequency with which the Sellers and the Buyers used the pooling equilibrium (as defined by our conservative definition) in the monetary and no-competition treatments.

Table 4: Individual Frequencies of Pooling Equilibrium Strategies

	Sellers' Behavior		Buyers' Behavior	
Treatment	first 5 blocks	last 5 blocks	first 5 blocks	last 5 blocks
Monetary	60%	61%	42%	55%
No Competition	5%	5%	25%	26%
Monetary vs No Competition	$p < 0.001$	$p < 0.001$	$p = 0.011$	$p < 0.001$

Notes: 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 the 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 treatment, where similar means cutoffs are no more than 0.05 away from each other. If we increase the band to 0.1, are 30% of Buyers play the pooling equilibrium strategy in the first half of the no-competition treatment, whereas 28% play in the second half.

As Table 4 shows, in the second half of the experiment, the majority of both Buyers and Sellers in the monetary treatments adhere to the pooling-equilibrium strategies far more than they do in the no-competition treatment. In the monetary treatment, more than 60% of Sellers send the same message irrespective of the quality of the product they own; as a result, 55% of Buyers ignore messages completely. The situation is very different in the no-competition treatment, in which less than 10% of Sellers play the pooling-equilibrium strategy, and less than a third of Buyers do so.

Overall, our results show the monetary treatment displays outcomes and behavior that are significantly closer to the pooling equilibrium than those in the no-competition treatment. This finding validates the idea that the presence of induced psychological payoffs shifts outcomes from the pooling equilibrium toward the more informative type of equilibrium. Moreover, we observe learning in the monetary treatment in the direction of the pooling equilibrium over the course of the experiment, whereas no such trend exists in the no-competition treatment. The fact that some of

[25] Note for the monetary treatment, this definition is very conservative for Sellers, because it only considers behavior consistent with a pooling equilibrium if the Seller sends the same message regardless of the quality of the good received, whereas we know that many other strategies might be consistent with pooling-equilibrium behavior, that is, any behavior whereby the messages sent are not sufficiently informative.

our results in the monetary treatment deviate from that 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.

Result 1: *Market outcomes and behavior of Sellers and Buyers in the monetary treatment are much closer to the pooling equilibrium than what we observe in the no-competition treatment. The majority of Sellers in the monetary treatment send uninformative messages, which the majority of Buyers ignore. This results in trade frequencies that are significantly lower in the monetary treatment than in the no-competition treatment. Overall, the introduction of psychological payoffs positively affects trade in the markets and shifts behavior toward some information transmission.*

4.2 Impact of Competition on Market Performance and Welfare

Now that we have isolated the impact of psychological payoffs on behavior, we turn to the impact of competition on market outcomes and welfare. Recall that, in theory, competition does not introduce any new equilibria. Thus, whether the competition is good or bad for the Buyers is a pure empirical question. If Sellers use the same strategies and Buyers interpret messages the same way in both games, the competition among Sellers should deliver higher payoffs for the Buyer because the chance is greater that at least one of the buyers will be endowed with a high-quality good. The same prediction would hold if a "better" equilibrium is played in the game with competition, for instance, PIE2 instead of PIE1. If, instead, subjects settle on PIE1 or the pooling equilibrium in the presence of competition, as opposed to PIE2 played in the game without competition, Buyers' welfare may decrease.[26] Our data show which of these assertions holds water in our experiment.

Figure 4 depicts the frequency of Buyers' purchasing decisions in each treatment as well as the quality of the purchased good.

As Figure 4 shows, competition has a significant effect on the purchasing decisions of Buyers, especially in the second half of the experiment (last 5 blocks), which arguably represents the most relevant data because they capture the behavior of subjects after they have had the time to learn the game and converge to a stable strategy. Figure 4 shows two main results. First, competition increases trade: over the last five blocks of the no-competition treatment, only 43% of the products available were actually sold, whereas this percentage increases to 57% when competition exists between Sellers. Second, Buyers end up with lower average quality of the product when competition is present. The statistical tests confirm these observations. In the last five blocks of the experiment, competition leads to Buyers purchasing products more often ($p < 0.01$) and to a lower average quality of purchased goods (significant at the 10% level, $p = 0.098$).[27]

Table 5 reports regressions that investigate the effect of competition on Buyers' and Sellers' welfare. To make comparison between Sellers' payoffs in the two treatments valid, 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 accumulated in the belief-guessing task and focus only on the tree-game payoffs.

[26] These are, of course, equilibrium considerations, assuming subjects are playing one of the characterized equilibria. Out of equilibrium, many things can happen that would affect ranking between treatments in terms of Buyers' welfare.

[27] These quantities are not statistically significant in the first five blocks: $p = 0.223$ and $p = 0.840$ in two comparisons, respectively.

Figure 4: Aggregate Outcomes, by Treatment

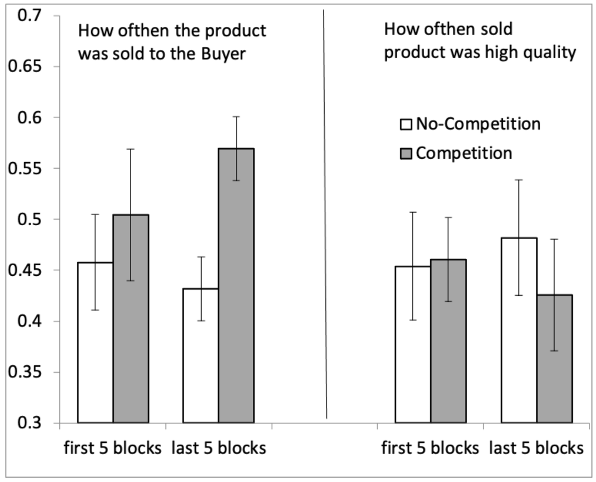

Notes: The left panel depicts purchasing frequency by treatment. The right panel depicts the likelihood that the product was high quality conditional on the product being purchased. Bars indicate 95% confidence intervals using robust standard errors, which are computed by clustering observations by session.

Table 5: Effect of Competition on Payoffs of Buyers and (selected) Sellers in the Tree Game

	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 accumulated for guessing beliefs tasks. Standard errors are clustered at the session level. ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

A few interesting patterns emerge from this analysis. 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. 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.

Which types of Sellers and Buyers suffer the most from the presence of competition? Table 6 presents average payoffs of Sellers and Buyers broken down by their types in each of the treatments in the last five blocks.[28]

Table 6: Which Types of Buyers and Sellers Suffer the Most from Competition (last 5 blocks)?

		No-Comp	Comp	Difference
SELLERS low quality product	type S1 ($G = 0, L = 0$)	12.20 (0.84)	14.93 (1.03)	YES** ($p = 0.05$)
	type S2 ($G = 6, L = 0$)	8.26 (0.54)	6.67 (0.76)	YES* ($p = 0.08$)
	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*** ($p < 0.01$)
SELLERS high quality product	type S1 ($G = 0, L = 0$)	7.46 (0.28)	7.83 (0.23)	NO ($p = 0.20$)
	type S2 ($G = 6, L = 0$)	7.69 (0.26)	8.52 (0.41)	NO ($p = 0.15$)
	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 ($p = 0.54$)
BUYERS	$\omega \leq 0.2$	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 ($p = 0.18$)
	$0.4 < \omega \leq 0.6$	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 (see Approach to Data Analysis for details). *, **, and *** indicates significance at the 10%, 5%, and 1% level, respectively.

As Table 6 shows Sellers who own the high-quality product earn the same average payoffs in the no-competition and competition treatments, irrespective of their psychological type. The Sellers with low-quality goods 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 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.

Results 2: *Competition encourages Buyers to purchase goods more often even though they end up with a low-quality product more often. Both Buyers and selected Sellers earn lower average payoffs in the presence of competition, with the largest payoff losses experienced by Sellers with low-quality goods and high sensitivity to lying and guilt and by Buyers with high sensitivity to disappointment.*

4.3 What Drives these Aggregate Results?

To explain the detrimental effects of competition on both Buyers' and Sellers' welfare, we proceed by investigating participants' strategies and beliefs to dissect the mechanism responsible for this result.

[28] In the Online Appendix, we present the same statistics for the first five blocks of the experiment.

Recall our experiment elicits detailed information about both Buyers' and Sellers' strategies and beliefs. In each game, with and without competition, we observe the stated beliefs and purchasing cutoffs of Buyers for both messages m_0 and m_1. Beliefs for message m_i indicates the likelihood that this message was sent by a high-quality Seller. The cutoff for message m_i, denoted by $\omega'(m_i)$,, defines the Buyers' purchasing decisions given her disappointment sensitivity; that is, types $\omega \leq \omega'(m_i)$ will purchase the product after observing m_i, whereas the remaining types will not. Additionally, in the competition treatment, a Buyer revealed her choice of Seller whom she would choose conditional on the combination of messages she could receive from the two Sellers. As for the Sellers, we observe their communication strategies in each treatment, that is, messages they want to send to the Buyer for all possible combinations of product quality they own and their psychological types (lying and guilt parameters), which results in eight possible scenarios. In addition, we observe Sellers' second-order beliefs regarding the first-order beliefs Buyers hold for each message.

4.3.1 Seller's Messages

Figure 5: Communication Decisions of Sellers (last 5 blocks)

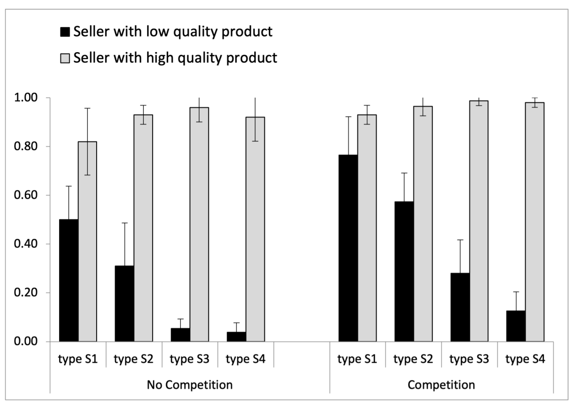

Notes: Average frequency of sending message m_1 is presented for each type of Seller in each treatment in the second half of the experiment. We compute 95% confidence intervals using robust standard errors obtained by clustering observations by session.

We first turn to Sellers' behavior. Figure 5 presents the average frequencies of m_1 messages sent by Sellers of different psychological types in the last five blocks of the experiment.[29]

The Sellers who own the high-quality product are expected to tell the truth and send message m_1 irrespective of their psychological type. Figure 5 shows this prediction is borne out in the vast majority of cases: over 80% of messages sent by high-quality Sellers in the second half of the experiment in both treatments are truthful. To perform statistical analysis, we estimate the 95% confidence interval around the observed average likelihood that Sellers with high-quality products

[29]The Sellers' communication decisions in the first five blocks look very similar and are presented in the Online Appendix.

send message m_1, and we look at whether our theoretical prediction (which says this likelihood should be 1) falls into this 95% confidence interval or not. In our no-competition treatment, the theoretical prediction is contained in the 95% confidence interval for psychological types S3 and S4 but not for the types S1 and S2. In the competition treatment, the theoretical prediction is contained in the 95% confidence interval for types S2, S3, and S4 but not for S1.

The situation changes when we look at Sellers with low-quality goods. Theoretically, in PIE 1, both types S1 and S2 are expected to lie and send message m_1, whereas in PIE2, only type S1 should lie. 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 about 60% of types S2), whereas types S3 and S4 lie much less. 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.[30] In other words, although Sellers' strategies do not conform to point predictions of either of the PIEs, they are in line with our key qualitative intuition of a PIE, which predicts that low-quality Sellers with larger psychological costs are less likely to send a potentially misleading message m_1.[31]

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 (see Figure 5). For example, in the no-competition treatment, whereas Seller 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, these percentages increased to 77%, 57%, 28%, and 13%, respectively, when 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).

An alternative way to support this claim is to use stacked regression analysis, where we analyze messages sent by the Sellers who own low-quality products, conditional on their psychological types and the treatment (see Table 7). These regressions confirm that Sellers with higher psychological costs from guilt and lying tend to lie less in the communication stage in both treatments. Moreover, the amount of lies observed between the two treatments increases significantly, conditional on the psychological type of a Seller: Sellers with low-quality products lie more in the competition than in the no-competition treatment.

Result 3: *In both games with and without competition, Sellers with high-quality goods tend to tell the truth, whereas Sellers with low-quality goods often lie. Those with higher values of guilt and lying sensitivity (types S3 and S4) lie less than those with lower values. Competition leads to more lying by Sellers who own low-quality products, irrespective of their psychological type.*

Table 7: Messages Sent by Sellers with Low-Quality Products

	first 5 blocks	last 5 blocks
Seller's psychological type		
S2	-0.86*** (0.20)	-0.47*** (0.25)
S3	-1.63*** (0.25)	-1.28*** (0.41)
S4	-1.94*** (0.38)	-1.79*** (0.51)
Indicator for Competition treatment	1.06*** (0.21)	1.29*** (0.23)
Constant	0.07 (0.23)	-0.31 (0.27)
# of observations	2178	2184
# of clusters	7	7
Log pseudo-likelihood	-1303.09	-1300.75

Notes: LOGIT regressions with a dependent variable equal to 1 if the Seller sent message m_1, and 0 otherwise. Omitted category is the Seller with a low-quality product and psychological type S1. Observations are clustered at the session level. *** indicates significance at 1% level.

Table 8: Buyers' Beliefs and Purchasing Cutoffs, Sellers' Beliefs, and Quality of Products for Different Messages

	$\omega'(m_i)$	$\bar{b}^1_B(m_i)$	$\bar{b}^2_S(m_i)$	$\bar{q}_H(m_i)$	$\bar{b}^1_B(m_i) = \bar{b}^2_S(m_i)$	$\bar{b}^1_B(m_i) = \bar{q}_H(m_i)$	$\bar{b}^2_S(m_i) = \bar{q}_H(m_i)$
first 5 blocks							
No Competition							
message m_0	0.31 (0.02)	0.24 (0.03)	0.22 (0.03)	0.07 (0.04)	$p = 0.218$	$p < 0.001$	$p < 0.001$
message m_1	0.55 (0.03)	0.71 (0.02)	0.73 (0.02)	0.66 (0.03)	$p = 0.561$	$p = 0.001$	$p < 0.001$
Competition							
message m_0	0.34 (0.02)	0.24 (0.06)	0.25 (0.02)	0.22 (0.02)	$p = 0.680$	$p = 0.757$	$p = 0.052$
message m_1	0.56 (0.04)	0.75 (0.04)	0.68 (0.02)	0.51 (0.03)	$p = 0.005$	$p < 0.001$	$p < 0.001$
last 5 blocks							
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 three columns are estimated from participants' strategies. Specifically, $\omega'(m_i)$ denotes the highest disappointment parameter for which the Buyer purchases the product after observing m_i, whereas $\bar{b}^1_B(m_i)$ and $\bar{b}^2_S(m_i)$ denote the 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 in the first four columns, 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 (six column), and Sellers' beliefs and the average actual frequency of high-quality Sellers for different messages (seventh column).

4.3.2 Buyer's Beliefs and Purchasing Decisions

The second and third columns in Table 8 present Buyers' beliefs and cutoffs conditional on receiving two messages, m_0 and m_1. Upon receiving message m_0 in both partially informative equilibria,

[30] In the last five blocks of the no-competition treatment, we observe significant difference between the frequency of sending m_1 message by types S1 and S2 ($p = 0.003$), and between S2 and S3 types ($p < 0.001$), whereas no

Buyers should set their belief to zero and not buy the good. In our data, Buyers' average beliefs during the last five blocks are 0.26 and 0.22, whereas their purchasing cutoffs were 0.29 and 0.31 in the no-competition and competition treatments, respectively. Because these are corner solutions, all the deviations are going to be positive and will move average Buyers' beliefs and corresponding cutoff away from the prediction of zero. Thus, unsurprisingly, we reject this prediction.[32]

A perhaps more informative statistic might be a simple descriptive statistic of how often reported beliefs and cutoffs were close to zero, allowing for some small noise. In both treatments, in the last five blocks, the majority of reported beliefs and purchasing cutoffs upon observing message m_0 are at most 5 percentage points away from zero (no-competition treatment: 63% for beliefs and 64% for cutoffs; competition treatment: 77% for beliefs and 74% for cutoffs).

The two PIEs that can be sustained in our games differ in Buyers' beliefs and purchasing cutoffs after receiving message m_1: PIE 1 is characterized by a lower belief $b(m_1) = 0.57$ and purchasing cutoff of $\omega^{\bar{\text{PIE2}}}(m_1) = 0.51$, whereas PIE2 is characterized by a higher belief $b(m_1) = 0.73$ and a corner solution for purchasing decision $\omega^{\bar{\text{PIE2}}}(m_1) = 1$. In the latter equilibrium, even the Buyer who is most sensitive to disappointment prefers to purchase the product, because only the Seller without any psychological costs and a low-quality product sends m_1 in equilibrium.

The beliefs of our Buyers, summarized in Table 8, favor PIE2 over PIE1. In fact, although we reject the PIE1 prediction for Buyers' beliefs in both games, these average beliefs are not significantly different from beliefs predicted by PIE2.[33] At the same time, Buyers' purchasing cutoffs upon receiving message m_1 are lower than those predicted by PIE2 but higher than those predicted by PIE1. Table 8 depicts this information. In the last five blocks, we observe an average cutoff of 0.59 and 0.62 in the game with and without competition, respectively, which is significantly higher than the one predicted by PIE1 but significantly lower than the one predicted by PIE2.[34]

The explanation for why our Buyers form beliefs consistent with the PIE2 prediction about the meaning of the message m_1 while simultaneously using a *lower* purchasing cutoff than the one predicted by PIE2 is simple. Recall that our theoretical analysis was performed under the assumption that Buyers are risk neutral. However, if one would allow Buyers to have concave utilities, that is, be risk averse, we could observe Buyers holding the "correct" beliefs but using lower cutoffs, which is precisely what our data indicate.

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 probability 0.50 or 0.40 for each point invested in the investment task 1 and 2, respectively. Thus, subjects who invest

---

significant difference exists between the frequency of sending m_1 for types S3 and S4. In the last five blocks of the competition treatment, all pairwise comparisons between the frequency of sending message m_1 of Sellers with low-quality products with different psychological types are statistically significant with $p < 0.001$.

[31] If anything, Sellers' behavior in the no-competition treatment is closer to PIE2 with a significant fraction of honest types who do not lie even absent induced psychological costs in our experiment. By contrast, Sellers' behavior in the competition treatment is closer to PIE1 rather than PIE2.

[32] Each of these averages is significantly different from zero in both treatments ($p < 0.001$).

[33] For the no-competition treatment, we obtain $p = 0.272$ and $p = 0.259$ for the beliefs in the first five and last five blocks of the experiment, respectively. For the competition treatment, we obtain $p = 0.634$ and $p = 0.236$ for the first five and last five blocks. At the same time, we reject the hypothesis that average Buyers' beliefs are equal to 0.57, which is the point prediction in PIE1 in all these cases ($p < 0.001$).

[34] We reject the null hypothesis that observed average cutoffs equal those predicted by PIE1 with $p = 0.09$ and $p = 0.01$ for the no-competition and competition treatment, respectively. We also reject that these average cutoffs equal 1, which is what PIE2 predicts, with $p < 0.001$ in both games.

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 ($p = 0.066$).[35]

Interestingly, Buyers' beliefs and purchasing cutoffs are remarkably similar in two treatments: with and without competition (see Table 8). Focusing on the last five blocks of the experiment, we note 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). Similarly, we find that Buyers use the same purchasing cutoffs in both treatments. Conditional on receiving an m_0 message, these cutoffs are 0.29 and 0.31 for the no-competition and competition treatment, respectively ($p = 0.627$), whereas upon receiving the m_1 message, these cutoffs are 0.59 and 0.62 for the no-competition and competition treatments, respectively ($p = 0.558$). In other words, Buyers did not pick up on the inflation of language that occurs when Sellers compete with each other and interpret messages in a similar manner in both games.

We now ask how a Buyer selects a Seller in the competition treatment. Two distinct scenarios are considered: one in which both Sellers send the same message, and one in which two Sellers send different messages. In the first scenario, the Buyer has no basis to prefer one Seller to the other, which is confirmed in our data. In the last five blocks of the experiment, Buyers chose Seller 1 in such a situation 53% of the time; this fraction is not statistically different from 50% ($p = 0.527$).[36] At the same time, when Buyers gets two different messages, they mostly chose the Seller with message m_1 rather than m_0, which happened in 84% of the cases in the last five blocks of the experiment, which is significantly different from 50% ($p < 0.001$).[37]

Result 4: *In both games with and without competition, Buyers rarely purchase the product upon receiving message, indicating the product is low quality (m_0), but do so often upon receiving the message that the product is high quality (m_1). The combination of Buyers' beliefs and purchasing cutoffs is consistent with Buyers being risk averse and playing the more informative PIE (PIE2 rather than PIE1). When Buyers have a choice between two Sellers, they select the Seller who sends message m_1 if two messages are different, and otherwise pick a Seller randomly. The presence of competition does not change Buyers' average beliefs and purchasing cutoffs.*

4.3.3 Sellers' Beliefs

The final piece of the puzzle is the second-order Sellers' beliefs about first-order Buyers' beliefs and its connection to true (observed) quantities. In any equilibrium, we expect to see these beliefs

[35] To reach this conclusion, we use the ORIV (Obviously Related Instrumental Variables) technique developed by Gillen, Snowberg, and Yariv (2018), which allows to correct for the measurement errors in elicitation of risk attitudes. ORIV is an improved version of 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.

[36] The same is true in the first five blocks of the experiment: upon observing identical messages, Buyers choose Seller 1 in 54% of the cases, which is not statistically different from 50% ($p = 0.233$).

[37] The same fraction is 80% in the first five blocks of the experiment, which is significantly different from 50% ($p < 0.001$).

coincide and confirmed:

$$\bar{b}_B^1(m_i) = \bar{b}_S^2(m_i) = \Pr\left[q = q_H | s^S(q,g,l) = m_i\right] = \bar{q}_H(m_i) \quad \forall m_i \in M.$$

Table 8 presents observed beliefs for each message m_i and the observed probability that message m_i comes from the high-quality Seller, $\bar{q_H}(m_i)$. This table also presents statistical tests comparing the three quantities. A few interesting patterns emerge. Focusing on the last five blocks in the experiment, we first note that Buyers' and Sellers' beliefs are similar to each other in both treatments. No statistical difference exists between beliefs for message m_1 in the no-competition treatment ($p = 0.136$) and message m_0 in the competition treatment ($p = 0.318$). For the remaining two cases, message m_0 in the no-competition treatment and message m_1 in the competition treatment, a statistical difference exists, but the average beliefs are very close to each other. Second, except for message m_0 in the competition treatment, our Buyers do quite poorly because they are consistently overestimating the likelihood that message m_i comes from the Seller who owns a high-quality product. The difference is especially pronounced in the competition treatment for message m_1 for which average Buyers' beliefs are almost 30 percentage points above actual ones.

Result 5: *Sellers understand how Buyers interpret messages, because their second-order beliefs either match or come very close to first-order beliefs of Buyers in both treatments. Buyers often interpret messages incorrectly. They correctly interpret message m_0 in the competition treatment and overestimate the proportion of high-quality Sellers who send message m_i in all other cases. This effect is the most severe for message m_1 in the competition treatment.*

4.4 Putting the Puzzle Together

The aggregate results (section 4.2) paint a clear picture of the effects of competition among Sellers. Buyers unequivocally suffer in the presence of competition, because they are more likely to purchase lower-quality products and consequently earn lower average payoffs. Sellers suffer from competition as well, because both those who are selected to trade with a Buyer and those who are not earn on average lower amounts than a Seller in the game without competition.

The analysis of participants' strategies and beliefs (section 4.3) uncovers the mechanism through which competition affects market outcomes. The decreased welfare in the competition treatment is the result of a number of factors. First, Buyers tend to use the same strategies in both treatments; that is, they interpret messages the same way and use similar purchasing cutoffs in the game with and without competition. At the same time, Sellers with low-quality products lie more often in the competition treatment in order to be selected to trade with Buyers. Second, Buyers on average hold correct beliefs about the interpretation of messages sent by Sellers in the no-competition treatment, but consistently overestimate the likelihood of message m_1 coming from the high-quality Sellers in the competition treatment. In other words, competition leads to inflation of messages' meaning, but Buyers fail to recognize such change. Third, although Sellers anticipate Buyers' failure to adjust for message inflation in the competition treatment, they cannot use this insight to their advantage. As a result, Sellers earn lower average profits in the presence of competition, because those chosen by the Buyers suffer from high psychological costs of misrepresenting the quality of the product they own. This is, however, the best Sellers can do given Buyers' beliefs and actions.

We conclude this section by looking at the informativeness of messages in the presence and absence of competition. According to the Observation 1 (section 2.6), the informativeness is defined as the difference in beliefs upon observing an m_1 and an m_0 message. In other words, this

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

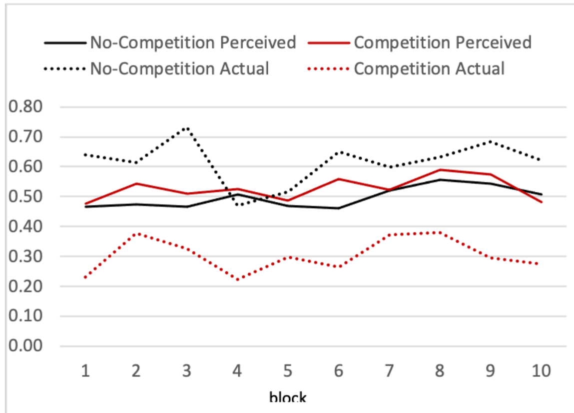

Notes: The solid lines depict the difference between Buyers' beliefs about message m_1 and message m_0 averaged in a block in each treatment. The dotted lines are the difference between probabilities that the product is high quality conditional on message m_1 versus m_0 for actual realized trades.

measure tells us how much more likely message m_1 is to be sent by the Seller with a high-quality product quality than message m_0. Figure 6 presents two measures of informativeness: the *perceived informativeness*, which captures Buyers' beliefs, and the *actual informativeness*, which depicts the correct market interpretation of messages given Sellers' behavior.

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.[38] 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.[39] 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.

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

[39]We perform the same statistical analysis for the competition treatment as the one reported in footnote 38 and obtain the following $p-$values for each block: $p < 0.001$, $p = 0.073$, $= 0.150$, $p = 0.004$, $p = 0.107$, $p < 0.001$, $p = 0.025$, $p < 0.001$, $p = 0.002$, $p = 0.014$.

5 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 buying subpar goods, whereas buyers 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. 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 evidenced by 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 psychologically 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.

With respect to competition, our results indicate 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 for their welfare, because they fail to change the way they interpret messages when competition is introduced. The abundant feedback and experience we provide out subjects does 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.

References

Attanasi, Giuseppe and Rosemarie Nagel. (2008) "A Survey of Psychological Games: Theoretical Findings and Experimental Evidence," in A. Innocenti and P. Sbriglia (Eds.) Games, Rationality and Behavior. Essays on Behavioral Game Theory and Experiments, Palgrave McMillan, 204-232.

Battigalli, Pierpaolo, and Martin Dufwenberg (2007). "Guilt in Games," *American Economic Review* 97(2): 170-176.

Battigalli, Pierpaolo , Dufwenberg, Martin, and Smith, Alec (2015). "Frustration and Anger in Games," *CESifo Working Paper Series* No. 5258.

Brandts, Jordi, and David Cooper (2007). "It's What You Say, Not What You Pay: An Experimental Study of Manager-Employee Relationships in Overcoming Coordination Failure," *Journal of the European Economic Association*, 5(6): 1223-1268.

Casella, Alessandra, Navin Kartik, Luis Sanchez, and Sebastien Turban (2018). “Communication in Context: Interpreting Promises in an Experiment on Competition and Trust,? *Proceedings of the National Academy of Sciences*, early edition.
Charness, Gary, and Martin Dufwenberg (2006).“Promises and Partnership,” *Econometrica* 74: 1579-1601.

Charness, Gary, and Martin Dufwenberg (2011). “Participation,” *American Economic Review* 101: 1213-39.

Crawford, Vincent, and Joel Sobel (1982). “Strategic Information Transmission,” *Econometrica*, 50(6): 1431-1451.

Dal Bo, Pedro, and Guillaume Frechette (2018). “On the Determinants of Cooperation in Infinitely Repeated Games: A Survey,” *Journal of Economic Literature* 56(1): 60-114.

Dufwenberg, Martin, and Uri Gneezy (2000). “Measuring Beliefs in an Experimental Lost Wallet Game,” *Games and Economic Behavior* 30: 163-182.

Dufwenberg, Martin, Maros Servatka, and Rado Vadovic (2017). “Honesty and Informal Agreements,” *Games and Economic Behavior* 102: 269-85.

Ellingsen, T., M. Johannesson, S. Tjotta, and G. Torsvik (2010). “Testing Guilt Aversion,” *Games and Economic Behavior* 68: 95-107.

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

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

Gibson R, Tanner C, Wagner A (2013). “Preferences for truthfulness: Heterogeneity among and within individuals,” *American Economic Review*, 103: 532-548.

Gillen, Ben, Snowberg, Erik, and Leeat Yariv (2019). “Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study,” *Journal of Political Economy*, 127(4): 1826-1863.

Goeree, Jacob, and Jingjing Zhang (2014). “Communication and competition,” *Experimental Economics* 17(3): 421-438.

Gneezy, Uri (2005). “Deception: The Role of Consequences,” *American Economic Review*, 95(1): 384-394.

Gneezy, Uri, Agne Kajackaite, and Joel Sobel (2018) “Lying Aversion and the Size of the Lie,” *American Economic Review* 108(2): 419-53.

Goeree, Jacob, and Leeat Yariv (2011). "An Experimental Study of Collective Deliberation," *Econometrica*, 79(3): 893-921.

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

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

Jin, Ginger, Luca, Michael, and Daniel Martin (2021). "Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure," *American Economic Review: Microeconomics* 13(2): 141-173.

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

Oprea, Ryan, Gary Charness, and Dan Friedman (2014). "Continuous Time and Communication in a Public-goods Experiment," *Journal of Economic Behavior and Organization*, 108: 212-223.

Schotter, Andrew, Weiss, Avi, and Inigo Zapater (1996). "Fairness and survival in ultimatum and dictatorship games," *Journal of Economic Behavior and Organization* 31(1): 37-56.

Smith, Vernon (1976)."Experimental Economics: Induced Value Theory", *American Economic Review*, 66 (2): 274-79.

Serra-Garcia, Marta, Eric van Damme, and Jan Potters (2013). "Lying About What You Know or About What You Do?," *Journal of the European Economic Association* 11 (5): 1204-1229.

Sobel, Joel (2020). "Lying and Deception in Games," *Journal of Political Economics* 128 (3): 907-947.

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

Vespa, Emanuel and Alistair Wilson (2016). "Communication with multiple senders: An experiment," *Quantitative Economics*, 7: 1-36.