Trust Me: Communication and Competition in Psychological Games*

Marina Agranov† Utteeyo Dasgupta‡ Andrew Schotter§

January 20, 2020

Abstract

We study theoretically and experimentally a communication market game with and without competition and embed it in a psychological game framework where sellers and buyers incur psychological costs. While the introduction of psychological costs is welfare increasing the introduction of competition is not. This is driven by inordinate amounts of lying by sellers when competition exists while buyers repeatedly tend to trust these dishonest sellers. Finally, in a contrasting treatment where subjects are not subject to psychological costs we find behavior falling back to the standard communication game principal agent problem with a babbling equilibrium without much transmission of information.

1 Introduction

Ever since the seminal paper of Crawford and Sobel (1982), economists have devoted a great amount of attention to communication games. In these games there is typically an informed principal who sends a message to a less informed agent who then takes an action that determines the payoffs of both people. When the preferences over material payoffs of the agent and the principal differ, the game takes on interesting strategic aspects. In its classic form neither the principal nor the agent has any aversion to lying, nor do they feel guilty when they mislead their partner, and are concerned only with maximizing their material payoffs from the game.

In many real-world communication games, however, there are psychological costs to lying and misleading others. For example, principals may feel guilty when they knowingly send a misleading message that will disappoint the agent. Consequently, if the principal experiences a high enough guilt for misleading, she may decide to refrain from lying. Additionally, some principals may simply have an aversion to deceiving an unsuspecting agent on principle. In such a situation the principal may choose not to lie even if that allows her to increase her material payoffs. When such psychological payoffs exist, an agent after receiving a message should not only consider the

*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, Wagner College University. Email: utteeyodasgupta@gmail.com.
§Department of Economics, New York University. Email: andrew.schotter@nyu.edu.

principal's preferences over material outcomes but also her attitude towards guilt and/or lies. Put differently the agent now needs to take into consideration not just the material preferences of the principal (in fact in our model they are known to the agent) but also her moral character. As a result if the agent knew that the principal is extremely averse to lying and/or guilt, then trusting her messages can become considerably easier for him. The problem however is that rarely does the agent know the moral character of a random principal. As a result, he has to strive to infer the actual type of the principal and the truthfulness of each received message.

The strategic situation we model in this paper is relevant for many different situations. For example, consider a politician (a mayor) trying to convince voters to support a proposition on the ballot which raises taxes to improve subway service in their city but also has a provision for allowing him to use those funds for other transportation needs he prefers (say roads). The voters are largely in favor of building subways but not roads. So in order to get the proposition passed the mayor will have to advocate subways even if he favors roads. If he comes out strongly in favor of subways the populace will have to assess how much of a bold face liar he is or how guilty he would feel if he told them he was in favor of subways, had them vote accordingly, and then used the money for roads. The issue is a matter of trust based on the voters' assessment of how sincere (averse to lying or averse to misleading voters) the mayor is. Further, if we introduce two politicians to the situation both bidding for the voters' attentions ("Crooked Hillary" and "Dissembling Donald Trump") the situation boils down to voters comparing the trustworthiness of the two candidates given their messages to the voters.

The examples above summarize the strategic setup we are interested in this paper. We look at principal-agent games where subjects have psychological payoffs that depend on the messages sent by principals and the beliefs about how much others are relying on their messages. In addition, we introduce competition between principals, where two principals send messages in an effort to win over the agent. The agent then decides which of the two principals to believe in. In particular, we consider a situation where two sellers (principals) send messages about the quality of their products to the buyer (agent), and the buyer has to decide on buying the product. We provide a model of such a "psychological-communication game," solve for its equilibria, both in the case of competition and no-competition, and then bring this model to the lab.

The psychological-communication game we study without competition involves two players: the seller and the buyer. The seller owns one unit of a product which 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, but not the low quality one. The situation is complicated by the fact that before purchasing the product the buyer cannot distinguish between the high and the low quality good and instead has to rely on the messages sent by the seller. In addition to the material payoffs, both players in the game obtain psychological payoffs as well, depending on their psychological types. Specifically, sellers might 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 may suffer from disappointment when their expectations about the quality of the good are misplaced. Our psychological-communication game with competition is very similar to the game without competition except for the introduction of two sellers who compete with each other to sell the good. The competition happens via communication, where both sellers send a message to a buyer who then decides to choose one of them based on the received messages.

Theoretically, the introduction of psychological payoffs into a communication game without competition is unambiguously beneficial since without them, the no-trade equilibrium is the only possible equilibrium outcome. In particular, the psychological-communication 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 in our psychological-communication game, however, is less clear. Although its introduction does not introduce any new equilibria in the market, it can affect which of the partially informative equilibria are selected. This is why we turn to controlled laboratory experiments for evidence on this issue.

In evaluating the results of our experiments we concentrate on the behavior of our buyers and sellers in markets with and without competition, and also look at the impact of the introduction of competition on the performance of these markets. Our aggregate data from the experiment reveal that in neither of the two games, one with competition and the other without competition, do the subjects play strategies consistent with a pooling equilibrium where sellers' messages are uninformative leading to no-trade outcomes. Instead, subjects play as if sellers' messages transmit information that can be relied upon by buyers. Moreover, in both games, observed buyer behavior is consistent with the most informative equilibrium predicted in our games, where only the sellers with no psychological disutilities would lie. We also find that in both games sellers do take into account their psychological costs and choose their messages strategically.

The above results however, are ancillary to our primary focus on analyzing the impact of competition between sellers in the psychological-communication game. To that end we find that the welfare of both sellers and buyers are unambiguously lower in the game with competition than in the game without competition. There are two effects that contribute to this decrease: (1) competition leads to higher levels of lying by sellers, even from those who have the greatest costs of doing so (i.e., those who suffer high lying costs and are highly guilt averse),[1] and (2) buyers fail to adequately adjust their beliefs to this acute dissembling from the sellers, indicating that they do not best-respond to the strategies that sellers play in the game with competition.[2] In fact the individual strategies of buyers seem to suggest that buyers' purchasing decisions are consistent with the belief that messages are *more informative* in the game with competition than without. So, buyers appear to be convinced that truth will more likely come out in the competitive environment, and as a result the competition environment encourages buyers to purchase goods more often even though they end up with a low quality product more often. Consequently, the average payoffs of both buyers and sellers are lower in the presence of competition.

Our paper has an interesting relationship to the extant literature on communication and psychological games.[3] The literature on cheap-talk communication in different types of strategic situations is extensive and is surveyed in Crawford (1998). More recent papers in this literature looked at effects of communication in the collective action settings (Goeree and Yariv (2011)), public good games (Serra-Garcia, van Damme, and Potters (2013) and Oprea, Charness, and Friedman (2015)), bargaining games (Agranov and Tergiman (2014) and Baranski and Kagel (2015)), coordination games (Cason, Sheremeta, and Zhang (2012), Brandts and Cooper (2007), Agranov and Schotter (2012)), and auctions environments (Agranov and Yariv (2017)). However, this paper is one of the very few papers that models a communication game as a psychological game. The seminal

[1]The idea here is that while in equilibrium none of these subjects should lie, the consequences of not being chosen for a sale by the buyer is such that their greed overcomes their natural distaste for lying and aversion to guilt.

[2]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.

[3]For a comprehensive survey of theoretical results in psychological games as of 2008, and existing experimental evidence see Attanasi and Nagel (2008).

experimental papers of Dufwenberg and Gneezy (2000) and Charness and Dufwenberg (2006) were only vaguely connected to communication games in the sense that subjects could communicate promises to each other but no formal equilibrium of the game was solved for. Their experiment focused on guilt aversion due to false promises, and found that promises, or statements of positive intent, improve trust, cooperation and efficiency. Gneezy (2005) studies the role of consequences in determining the level of deception. He finds that people care about both their own gain from lying and how much harm their lying may cause to the other side. Vanberg (2008) presented an experiment without an explicit theoretical model of psychological payoffs to distinguish between psychological motives such as promise induced changes of second order beliefs (where people dislike letting others down) vs. a preference for promise keeping under communication, and found that subjects had a preference for promise keeping on its own. See also Gneezy, Kahackaite, and Sobel (2018) for a theory and experimental evidence concerning lying in the individual decision-making task.

The experimental literature concerning the interplay between competition and communication is still in its infancy.[4] 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 experimentally communication game with hidden actions and communication among competing senders, but fail to model the game as a psychological game. The authors find that the meaning of messages depends on whether competition is present or not: messages are inflated in the game with competition, while these inflated messages induce mostly the same actions on the part of receivers.[5] Our game is the game with hidden information rather than hidden actions and our results show 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 they tend to believe these messages and act accordingly.

Closer to our setup, Goeree and Zhang (2014) present a game with similar material payoffs as ours and also introduce competition.[6] They find that 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. While 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 introducing psychological payoffs and modeling the game as a psychological game which allows players to exhibit a wide range of emotions (translated into their payoffs) such as aversions to guilt, lying, and disappointment. We then obtain theoretical results with regard to the effects of communication on market outcomes and players' behavior and test these predictions in a lab experiment where we induce and control the payoffs associated with these emotions. Despite different approaches, Goeree and Zhang (2014) and we, both show that competition tends to decrease efficiency when introduced in a communication game.

Finally, Born (2018) studies promise competition between sellers who differ in their intrinsic

[4]There are several studies that look at the effects of competition on trust in various environments (Huck et al. (2012), Keck and Krelaia (2012), and Fischbacher et al. (2009)).

[5]See also Gibson et al. (2013) who show that truthfulness is context dependent.

[6]For a communication game with hidden information but without competition see Charness and Dufwenberg (2011).

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 the welfare of buyers in comparison to a no-competition case because some sellers promise more than they would in the absence of competition. Experimental results reveal that behavior of sellers crucially depends on their experience in the game as the difference between competition and 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 at the end of the experiment, after subjects have learned to play the game and converged to a stable behavior.

In this paper we will proceed as follows. In Section 2 we introduce the model of the buyer seller communication game with psychological payoffs. This serves as the basis for the experiment we run. In Section 3 we describe our experimental design. In Section 4 we describe subject behavior in our experiment and the extent to which it aligns with the predictions of our model. Section 5 follows up on the discussions introduced in Section 4 to discuss the impact of competition on the welfare of our Buyer and Seller. Section 6 strips away the psychological payoffs we impose in our other treatments while Section 7 looks at the data at the individual level. Finally, Section 8 offers some conclusions and a discussion.

2 The Model

We study a straightforward communication game between an informed seller and an uniformed buyer. The informed seller wants to sell a good which can either be of high quality or of low quality. The seller gets to send a message about the quality of the product to the buyer. The seller can send one of the two possible messages: m_1= "The product is really high quality" and m_0 ="The product is low quality". The uninformed buyer does not know the quality of the good, but observes the message sent by the seller. Based on this message the buyer either buys, or not, and receives a payoff that depends on the good's quality. 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 that such emotions exist and must take them into account when evaluating the veracity of the messages he or she sends and receives.

We study this game in two contexts: the no-competition case where there is only one seller and one buyer, and the competition case, where there are 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 will first describe our model in the absence of psychological payoffs, i.e., using only material payoffs, and then introduce both psychological payoffs and competition.

2.1 Material Payoffs

There are two players: the Seller (he) and the Buyer (she). The Seller owns one unit of the product and wants to sell it to the Buyer. The quality of the product is low $q = q_L$ with probability $p > \frac{1}{2}$ and high $q = q_H$ with remaining probability $1 - p$. The Seller knows the quality of the product he owns. The Buyer is interested in purchasing the high quality product, but not the low quality product. The situation is complicated by the fact that the Buyer can not distinguish the high from the low quality good till she purchases it and has to instead rely on the messages sent by the Seller prior to making purchasing decisions.

At the outset of the game, Nature determines the quality of the good owned by the Seller based on the distribution specified above. The Seller observes the quality of his good and sends a message to the Buyer in an attempt to convince her to purchase his good. We will focus on the situation in which there are exactly two messages m_1 and m_0 (as defined above) and will denote the set of messages by $M = \{m_0, m_1\}$. After observing the Seller's message, the Buyer makes a decision whether to buy the product or not, after which the game ends. The material payoffs of players are depicted in Figure 1, where the top payoff at each end node is the Buyer's payoff and the bottom one is the Seller's payoff. The dashed line indicates the information set of the Buyer since she does not know the type of the Seller she is dealing with.

Figure 1: Material Payoffs in the Game without Competition

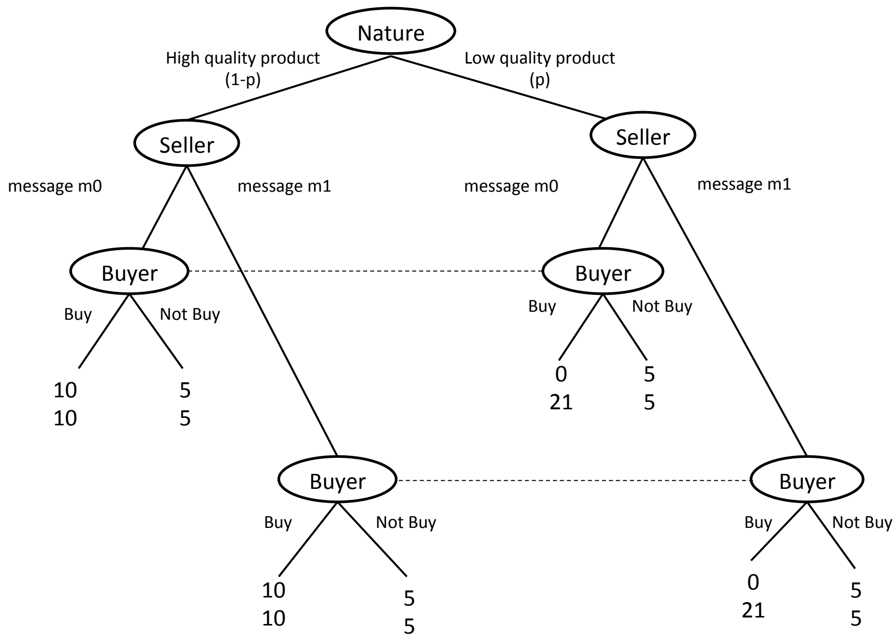

Notes: The top payoff at each node is the Buyer's material payoff and the bottom one is the Seller's material payoff.

As we see, when a good is not sold the Buyer and the Seller each receives 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. Since in this case the preferences of the Buyer and the Seller are completely 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 this is the case, assume by contradiction that there exists an equilibrium, in which after observing message $m_i \in M$ the Buyer purchases the product with a higher probability than after observing a message $m_j \in M$. Such behavior is justified if the Buyer believes that 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. This implies that there can be no equilibrium where one message entails a higher probability of the high quality product than another. Since the messages are uninformative, the Buyer is left with her original prior probability and since the prior likelihood that the product is of high quality is $1-p < \frac{1}{2}$, it follows that given our material payoffs, *the no-trade equilibrium is the only one that exists.*

2.2 Psychological Payoffs

The above analysis changes substantially when we introduce psychological payoffs in this game. 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 the important ones.[7]

Specifically, we assume that the Seller may experience two emotions, guilt and lying-aversion, while the Buyer may experience disappointment. Guilt stems from the fact that a Seller can feel badly 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 this negative emotion on Seller's payoff depends, of course, on the beliefs of the players. In contrast, a Seller may simply experience discomfort from lying whenever he knowingly sends a false message. In this case the negative emotion neither depends on how the Buyer interprets the message, nor does it depend on 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, define the Seller's type as a triple (q, G, L), where the first entry is the product quality $q \in \{q_L, q_H\}$, and the other two are psychological parameters indicating guilt $G \in \mathbb{G}$ and lying sensitivity $L \in \mathbb{L}$. The Seller's type is drawn from a joint distribution $F(\cdot)$, and the set of Seller's types is denoted by T^{Seller}. The Buyer's type is determined by a single parameter $\omega \in \Omega$ which captures Buyer's disappointment sensitivity. This parameter is drawn from the distribution $H(\cdot)$, and the set of Buyer's types is denoted by T^{Buyer}.

The Seller's strategy is a function that maps Seller's type into messages, i.e.,

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

where $s^S(q, G, L)$ denotes the message that the Seller with type (q, G, L) sends in the communication stage. The Buyer's strategy is a function that determines her purchasing probability given her type and the message received from the Seller, i.e.,

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

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 material payoffs described above, along with the psychological payoffs determined by players' beliefs and the messages sent by a Seller. We denote by $b_B^1(m_i)$ the first-order belief of the Buyer that the Seller has high quality product given message

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

m_i, and by $b_S^2(m_i)$ 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

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

To define the extent of disappointment, guilt, and lying, we use the fact that there are only two possible messages, and that they have natural meaning and a simple interpretation.[8] Recall that the two messages are $m_0 =$"The product is low quality" and $m_1 =$"The product is really of 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, i.e., the Seller dislikes lying. In our game lying comes up in two instances: if a Seller with a low quality good sends message m_1, then he is lying; similarly, if a Seller with a high quality good sends message m_0, then also he is 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.[9]

In 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 that 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 that the Seller experiences will depend on how sensitive the Buyer is to such disappointments. Other things equal, the higher the disappointment parameter ω of the Buyer, the more guilty the Seller feels when he leads the Buyer on and then sells her the low quality product. Formally, the amount of guilt that the Seller experiences is equal to

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

where $10 \cdot b_S^2(m_1)$ represents the Seller's belief regarding 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 will denote the last term $\mathbb{E}[\omega|\text{Buyer with type } \omega \text{ buys the product after receiving } m_1]$ 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 product which turns out to be of low quality. The reason we introduce Buyers' disappointment into the model is that without this emotion, it is unclear why Sellers should feel guilty for leading the Buyers on and then double crossing them. In other words, the reason Sellers feel guilty about convincing Buyers to buy the product is that Sellers know that Buyers would suffer disappointment if they do doublecross them. We define this disappointment to be equal to the difference between the expected and the actual material payoff the Buyer receives conditional on observing m_1, i.e.,

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

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

[9] In the experiment, we will abstract away from these last masochistic preferences, i.e., the situation in which Seller says m_0 when he has a high quality product. This simplification does not change the qualitative results we derive here.

This amount enters the utility of the Buyer multiplied by the disappointment sensitivity parameter ω. Collecting all the terms, we summarize players' payoffs as follows:

$$\Pi^{\text{Buyer}}(m_i, \text{Not Buy}) = 5 \qquad \text{for } m_i \in \{m_0, m_1\}$$

$$\Pi^{\text{Buyer}}(m_i, \text{Buy}) = \begin{bmatrix} 10 & \text{if } q = q_H \\ -\,10 \cdot b_B^1(m_i) \cdot \omega & \text{if } q = q_L \end{bmatrix} \qquad \text{for } m_i \in \{m_0, m_1\}$$

$$\Pi^{\text{Seller}}\,(m_i, \text{Not Buy}) = \begin{bmatrix} 5 - L \cdot 1_{m_i=m_0} & \text{if } q = q_H \\ 5 - L \cdot 1_{m_i=m_1} & \text{if } q = q_L \end{bmatrix}$$

$$\Pi^{\text{Seller}}\,(m_i, \text{Buy}) = \begin{bmatrix} 10 - L \cdot 1_{m_i=m_0} & \text{if } q = q_H \\ 21 - 10G \cdot b_S^2(m_1) \cdot \mathbb{E}[\omega|\text{ Buy}|m_1] \cdot 1_{m_i=m_1} - L \cdot 1_{m_i=m_1} & \text{if } q = q_L \end{bmatrix}$$

where $1_{m_i=m_1}$ ($1_{m_i=m_0}$) denotes the indicator function that takes the value one if Seller sent message $m_i = m_1$ ($m_i = m_0$) and zero otherwise.

Figure 2 incorporates the psychological payoffs discussed above into the game with material payoffs described earlier in section 2.1. We find in Figure 2 that psychological payoffs appear whenever the Seller lies and misleads the Buyer (Buyer's payoffs placed on top of Seller's payoffs). 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 that 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 $0 - 10 \cdot b_B^1(m_1)$ 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 $21 - (G \cdot 10 \cdot b_S^2(m_1) \cdot \omega) - L$. Here he receives a material payoff of 21 since he managed to peddle off a low quality good by claiming it to be a high quality one. However, since this leads to a disappointment of the Buyer of the amount $(10 \cdot b_B^1(m_1) \cdot \omega)$, 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, since he lied.[10]

2.3 Equilibrium Analysis of The Game without Competition

Equilibrium consists of specifying a communication strategy for the Seller, s^S, indicating the probability distribution over messages for each Seller's type, a buying strategy for the Buyer, s^B, indicating the probability that the Buyer buys the product for each message (m_0, m_1), and the system of beliefs of both the Buyer and the Seller (b_B^1, b_S^2) such that

(1) **Buyer's actions are optimal**

$$s^{B^*}(\omega, m_i) = \arg \max_{s^B \in [0,1]} \mathbb{E}\Pi^{\text{Buyer}}\left(m_i, s^B\right) \qquad \forall (\omega, m_i) \in T^{\text{Buyer}} \times M$$

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

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

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

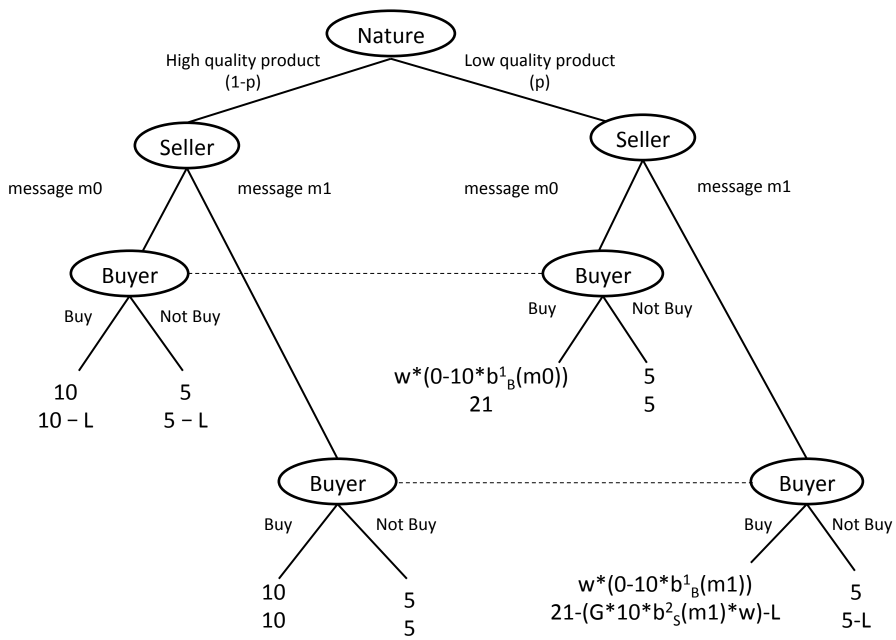

Notes: The top payoff at each node is the Buyer's overall payoff, while the bottom one is the Seller's overall payoff.

(2) **Seller's messages are optimal:**

$$s^{S^*}(q,G,L) = \arg\max_{m_i \in M} \mathbb{E}\Pi^{\text{Seller}}\left(m_i, s^{B^*}\right) \qquad \forall (q,G,L) \in T^{\text{Seller}}$$

(3) **Beliefs are correct:**

$$b^1_B(m_i) = b^2_S(m_i) = \Pr[q = q_H | s^S(q,G,L) = m_i] \qquad \forall m_i \in M$$

In words, just like in any psychological game, in equilibrium (a) actions of both players maximize their expected payoffs conditional on beliefs they hold regarding actions of other players, (b) no Seller type wants to mimic another type in terms of communication strategy used, i.e., messages are optimal, and (c) beliefs are 'correct', i.e., first-order and second-order beliefs of players coincide with the expected frequency of Buyer choosing to purchase the product conditional on received messages.

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

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

As standard in any cheap talk game, there exists a non-informative babbling equilibrium in our setting as well, in which messages sent by the Seller convey no information, and, therefore, the Buyer chooses the same action regardless of the message. We refer to this equilibrium as the pooling equilibrium and note that in our setting there exists a **unique pooling equilibrium**, 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, there exists **partially informative equilibria (PIE)**, 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 ones.

Specifically, in any PIE, for a given message m_i and belief $b_B^1(m_i)$ associated with it, the Buyer 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 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 (G, L), since (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 since $\frac{\partial \bar{\omega}(m_i)}{\partial b_B^1(m_i)} > 0$. Thus, message m_0 necessarily comes from the Seller with a low quality product, which implies that $b_1^B(m_0) = b_2^S(m_0) = 0$ and the Buyer does not to 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 the parameters of the game. Those with relatively lower values of guilt and lying sensitivities may lie and announce that they have high quality product in the communication stage, while 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, the Seller with low quality product and type (G, L) prefers to be truthful and send message m_0 if and only if

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

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

Finally, in equilibrium, beliefs must be correct, i.e.,

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

where ψ is the proportion of Sellers with a low quality product who send message m_1 in equilibrium. Note that the necessary condition for existence of PIE is that the proportion of Sellers with a low quality product who lie in equilibrium is not too high. This is required to guarantee that at least

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

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 in the game discussed in Section 2. Two sellers (Seller 1 and Seller 2) compete for the opportunity to sell the good to a single Buyer.

There are three players in this game now: the Buyer, Seller 1 and Seller 2. At the outset of the game, nature draws types of players: $\omega \in T^{\text{Buyer}}$ for the Buyer, $(q_1, G_1, L_1) \in T^{\text{Seller}}$ for Seller 1, and $(q_2, G_2, L_2) \in T^{\text{Seller}}$ for Seller 2. Player types are private information and are drawn independently from the same distributions as in the game without competition. After observing their types and the quality of their goods (each of which has a probability p to be 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.

In what follows, we discuss the properties of the equilibria that can be supported in the game with competition in intuitive terms, and refer the reader to the Appendix for a more rigorous discussion. Our analysis focuses on symmetric equilibria, in which both Sellers use the same communication strategy when they are of the same type, and hold the same beliefs.

The **symmetric equilibrium** in this game consists of specifying the communication strategy for the Sellers s^S, indicating the probability distribution over messages for each Seller's type, the selection function for the Buyer that indicates which Seller is selected by the Buyer based on observed messages, a buying strategy for the Buyer s^B, indicating the probability that the Buyer buys the product for each message (m_0, m_1) received from the selected Seller, and 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 communication strategy 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 communication strategy of the Sellers is optimal given their beliefs and the actions of the Buyer, and, finally, (c) all beliefs are correct.

In general, the game with competition admits the same *types* of equilibria as the game without competition. Specifically, there always exists a **unique pooling no-trade equilibrium** 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, while the unmatched Seller gets a payoff of zero. In addition, one can also support a **partially informative equilibria (PIE)** in our game with competition. In this case the messages are somewhat informative in the sense that Sellers that 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, then 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, then she randomly selects one of the Sellers and does not purchase the product. Finally, if the Buyer receives two m_1 messages, then 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 the game without competition and in the game with competition depends on the parameterization of the game. We discuss our chosen parameters in the next section which are also the same parameters we use in our experiment. These parameters were chosen so that the game with competition and the game without competition both have the exact 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 as it allows 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.

While we establish the existence of an equilibrium for our model using the parameters employed in our experiment, in a supplemental 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

In order to bring our model to the lab we need to set parameters to be used in both the competition and the no-competition games. To do this we will need to define a pair of probability distributions regarding types of buyers and sellers. Specifically, we assume that the probability that the product quality is low is $p = 60\%$, which ensures that absent of 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 $H(\omega) = U[0,1]$. Moreover, conditional on the product quality, there are four equally likely psychological types of Sellers: $(L, G) \in \{(0,0),(0,g),(l,0),(l,g)\}$ where $l = 20$ and $g = 6$.[13] This reflects the fact that in our experiment some Sellers are motivated by feelings of both guilt and lie aversion, some experience only one of these two emotions, 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 one that is 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 $\Pi^{\text{Seller}}(m_0, \text{Not Buy}) = 5$.

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, there exists three equilibria:

1. **Pooling equilibrium.** In this equilibrium, the Buyer treats messages from the Seller(s) as uninformative and does not update her prior beliefs about Seller's quality regardless of the observed message. That is, after observing either message, the Buyer believes there is 60% chance that she is facing a low quality Seller. In the game with competition, the Buyer randomly selects one Seller. In both games, the Buyer does not purchase the product and collects payoff of 5.

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

2. **PIE1.** In this equilibrium, Sellers with a low-quality product and with psychological types $(0, 0)$ and $(0, g)$ lie and send message m_1 in equilibrium, while 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, then 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, then the Buyer knows for sure that this message is sent by the low quality Seller, and thus, the Buyer 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 $(0, 0)$, i.e., those who do not suffer from either lying or guilt, lie and send m_1 in equilibrium. The remaining types truthfully reveal their product quality. In the game with competition, the Buyer selects a Seller with message m_1 if she receives two different messages, otherwise she randomly selects one Seller. If the chosen Seller's message is m_1, then the Buyer believes that there is 73% chance that this message comes from the high quality Seller and this belief is high enough so 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, then the Buyer knows for sure that the good is of low quality and thus does not buy the product. The Buyer's expected payoff is 6.04 in the game without competition, and 6.46 in the game with competition.

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

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

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 expected payoffs of the Buyer. The more information the Buyer receives from the Sellers' messages, the better purchasing decisions the Buyer 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, i.e.,*

$$\mathbb{E}\Pi_{POOL}^{Buyer} < \mathbb{E}\Pi_{PIE1}^{Buyer} < \mathbb{E}\Pi_{PIE2}^{Buyer}$$

2.6 Effect of Competition on Buyer's Welfare

As our previous discussion shows, our setup admits multiple equilibria in both games, with and without competition. Thus, 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 since different equilibria result in different expected Buyer's payoffs. This is ultimately an empirical question, and a compelling reason for conducting 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, then the Buyer's expected payoff is weakly higher in the game with competition.*

The logic behind Observation 3 is as follows: when messages are not informative, as is the case in the pooling equilibrium, there is no scope for competition to influence Buyers' payoffs. The benefit of competition between Sellers comes into play only when messages are partially informative as in the PIE. If the same PIE is played in both games, it means that Sellers' communication strategies 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 beneficial to the Buyer as it increases the chances of the Buyer playing the game with the high quality Seller.

However, there are a number of potential reasons why the Buyers' payoffs might decrease with competition. First, it is plausible that a Seller might be compelled to lie more often facing competition when he has a low quality product. Indeed, if he reveals the quality truthfully, then there is a greater chance that the other Seller will make the sale instead, 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!).[14] Thus, in the face of competition, we may see more lying and hence more opportunities for Buyers to purchase low quality goods. Second, it is possible that the introduction of competition will 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 for the Buyers or detrimental.

[14] Schotter, Weiss and Zapater (1996) have shown that competition can have an impact on rejection behavior in Ultimatum games where Receivers are more willing to accept low offers if the person making such an offer had to compete in a tournament-like setting.

2.7 Predictions

Given our parametrization, we have distinct point predictions about equilibrium behavior of Buyers and Sellers in our experiment. Since our model provides a multiplicity of equilibria, one of our goals in this paper is to try and detect which of these equilibria, if any, were adhered to by our subjects. However, since the proposed model is complex it is prudent to look for qualitative support for our theory and not just attempt to verify the point predictions. Consequently, at various points we examine weaker predictions that make ordinal rather than cardinal predictions.

Probably the easiest equilibrium to detect is the Pooling equilibrium since it makes the starkest predictions. In a Pooling equilibrium, no purchases should be made and Buyers' beliefs and actions should be the same irrespective of the messages received. This is true because the messages contain no information regarding the quality of the product that the Seller owns. In order for this to be true, however, there must be no correlation between product quality and messages sent by Sellers. This leads to Prediction 1.

Prediction 1: Behavior in Pooling equilibrium. In a Pooling equilibrium, Buyers' beliefs and actions are the same irrespective of the observed message of the (chosen) Seller, i.e., $b_1^B(m_1) = b_1^B(m_0)$ and $s^B(m_1|\omega) = s^B(m_0|\omega)$. In addition, since all messages are to be ignored, there should be no correlation between messages and product quality for the Sellers.

Our next prediction concerns the way Buyers interpret their received messages by converting them into beliefs about the product's quality. By looking at these beliefs we hope to gain some insight into the choice of equilibrium by our subjects.

Prediction 2: Message Interpretation in PIEs. In our experiment, $\bar{b}_B^1(m_1) = \{57\%, 73\%\}$ and $\bar{b}_B^1(m_0) = \{0\%, 0\%\}$, respectively, for our PIE1 and PIE2 equilibria, where $\bar{b}_B^1(m_i)$ is the mean Buyer beliefs in the experiment conditional on receiving message $m_i \in M$.

Prediction 2 provides the point predictions for mean Buyer beliefs for each of the two possible PIEs that might exist. More loosely, conditional on receiving message m_1, if Buyers' beliefs are close to 57%, then Buyers' beliefs adhere to predictions of PIE1, while if these beliefs are closer to 73% then Buyers' beliefs adhere to predictions of PIE2. In addition to the way in which Buyers interpret messages, our equilibria also provides a precise prediction of Buyers' purchase decisions conditional on received messages. This leads us to the next two predictions.

Prediction 3: Buyer Decisions in PIEs. In our experiment, $\bar{\omega}(m_1) = \{0.51, 1.00\}$ respectively for PIE1 and PIE2 equilibria, while in both PIEs $\bar{\omega}(m_0) = 0$, where $\bar{\omega}(m_i)$ is the mean sensitivity threshold set by the Buyer conditional on the received message $m_i \in M$.

Alternatively, evidence for a weaker form of Prediction 3 might only require that Buyers' cutoffs have the correct ordinal relationship as stated below.

Prediction 3a: Weak Buyer Decisions in PIEs. In any partially informative equilibrium, Buyers assign different probabilities to the events that messages m_1 and m_0 come from the Seller with a high quality product, i.e., $b_1^B(m_1) > b_1^B(m_0)$ and act accordingly, i.e., $s^B(\omega, m_1) \geq s^B(\omega, m_0)$.

Prediction 4: Sellers' Behavior in PIEs. In a PIE1, Sellers with low quality product and psychological types $(0, 0)$ and $(0, G)$ lie and send m_1, while in a PIE2, the only Seller who lies is the Seller with low quality product and psychological type $(0, 0)$. In either a PIE1 or a PIE2 equilibrium, no Seller lies when endowed with a high quality good.

The weaker version of the last hypothesis suggests that in any PIE there should be a monotonic relationship between messages and psychological types of Sellers conditional on having a low-quality product, while high-quality Sellers should always tell the truth.

Prediction 4a: Weak Sellers' Behavior in PIEs. In any PIE, Sellers with high quality products send message m_1, while Sellers with low-quality products satisfy

$$\Pr[s^S(q_L,0,0)=m_1] \geq \Pr[s^S(q_L,0,G)=m_1] \geq \Pr[s^S(q_L,L,0)=m_1] \geq \Pr[s^S(q_L,L,G)=m_1]$$

Furthermore, in any equilibrium all first and second order beliefs must be consistent and correct.

Prediction 5: Belief Confirmation. In all equilibria,

$$\bar{b}_B^1(m_i)=\bar{b}_S^2(m_i)=\Pr\left[q=q_H|s^S(q,L,G)=m_i\right] \text{ for all } m_i \in M$$

One interesting feature of our model and the associated experiment is the fact that even though the introduction of competition into our communication game does not introduce any new equilibria, and we still have only our Pooling, PIE1 and PIE2, competition does have consequences for the welfare of our Buyers and Sellers. This leads us to our next prediction.

Prediction 6: Buyers' Welfare and Competition. In any PIE, the payoffs of the Buyers are higher when there is a competition between Sellers.

Note that Prediction 6 compares the payoffs of Buyers when we hold the equilibrium fixed and introduce competition. However, it is possible that the introduction of competition may lead to a change in the equilibrium selected by the subjects and hence may have an impact on welfare via an equilibrium selection channel. To test for such a change we would need to investigate whether Buyer and or Seller behavior changed when competition was introduced. To simplify our prediction we merely state the following null hypothesis summarized as Prediction 7.

Prediction 7: Equilibrium Selection and Competition. There is no difference in the strategies of Buyers and Sellers between the games with and without competition.

3 Experimental Design

The experiment was conducted in the experimental lab of the Center for Experimental Social Science (CESS) at New York University. 179 subjects were recruited via E-mail from the general undergraduate population at NYU for an experiment that lasted approximately one hour and forty five minutes. Subjects received a show up fee of \$7 and on average received a final payment of \$29.5 for their participation. The program used in the experiment was written in Z-Tree (Fischbacher (2007)).

The experiment was a direct implementation of the model described above. We conducted three separate treatments: a No Competition treatment, a Competition treatment and a Monetary treatment. While the first two treatments incorporate psychological payoffs, the last one does not, and essentially is the situation depticted 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 they lie to the Buyer and disappoint her. We also impose a sensitivity parameter ω on the Buyer that specifies her sensitivity about being misled by the Seller. Note 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 where such costs are typically inferred. It is useful to point out 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 then 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.[15,16]

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. There is a 40% chance that the product is of high quality and 60% chance that it is of low quality. 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 or not based on the message she receives.

Consistent with the theoretical framework analyzed earlier, the Seller can be one of the four types with equal probability in the experiment: Type S1 has $(L = 0, G = 0)$, Type S2 has $(L = 0, G = 6)$, Type S3 has $(L = 20, G = 0)$, and Type S4 has $(L = 20, G = 6)$. Further, the disappointment parameter ω is drawn randomly and uniformly over the interval $[0, 1]$ for each Buyer. The Seller's task is to specify a strategy that is a function of his psychological type and the type of good he is endowed with to the set of messages $\{m_{0,}m_1\}$. The Sellers enter their strategies by filling out a table which 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 buying strategy that decides whether she buys the good conditional on the message she receives and her sensitivity type (ω). Buyers do that in the experiment by entering two cutoff values, $\omega'(m_0)$ for message m_0, and $\omega'(m_1)$ for message m_1, such that whenever the realized value of ω is less than $\omega'(m_0)$ $(\omega'(m_1))$ the Buyer buys the good (does not buy the good). This strategy 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. Specifically, each Buyer was asked to enter a number between 0 and 100 representing her belief that the Sellers that sent message m_i possessed a high quality good. This was done for both messages m_0 and m_1. Further, from the sellers we elicited their second order beliefs that specified what the Seller thought the first order belief of the Buyer was, upon receiving either message m_0 or m_1. The subjects were rewarded for their guesses according to a quadratic scoring rule. In the instructions

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

[16]One may worry that subjects participating in our experiments experience additional psychological feelings which are inherent part of their personality, i.e., in addition to penalties imposed by our payoff's structure subjects experience more guilt, lying aversion, and disappointment. If that was the case, then the Sellers should engage in less lying and the Buyers should purchase the product less often compared with what equilibrium analysis predicts. As we will see in our results section, this is the opposite from what we observe in our experiment, which minimizes the above mentioned concern.

Figure 3: The interface of Seller's Task

the subjects were strongly urged to report truthfully since we told them that they would maximize their payoff in the experiment by doing so. We did this because we were not interested what the best belief elicitation rule was and merely wanted subjects to report truthfully (for similar approach see Nyarko and Schotter (2002)).

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.[17]

We use this block design because entering a strategy and a set of beliefs and reviewing feedback is a time consuming process. Hence it would be very time consuming to have subjects do this for say 50 periods. Our design allows subjects to maximize the amount of feedback they get while economizing on the time they spend mechanically entering their strategies and beliefs. More importantly, we feel this is the correct way to conduct experiments using the strategy method since once a strategy is entered one might as well receive a lot of feedback on it before being asked to change it.[18] Entering a strategy and receiving only one period of feedback hardly allows a subject

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

[18]A similar method is used by Dal Bo and Frechette (2016) when they study infinitely repeated prisoners' dilemma

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. We did this because in psychological games, payoffs are a function of beliefs and it is very hard to isolate the pure incentive for truthful belief revelation. 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 so as 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 their guilt costs no matter what level his G turns out to be. Paying for either beliefs or game payoffs minimizes or eliminates this incentive. In our Supplementary Appendix we provide some calculations that suggest that given our payoffs, there is very little, if any, advantage to misrepresenting one's beliefs.

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 $R = 2.5$, while in the second task $p = 0.4$ and $R = 3$. One of these two risk tasks was randomly chosen to count 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 will buy from given the messages received from each. There were four different scenarios that 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 while 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** while 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 will 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 that 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.

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

4 Results

This section is organized as follows. First, we describe our approach to the data analysis and the general method we use to perform the statistical tests. Next we investigate the predictions made by our theory in Section 2.7 about our game with and without competition when psychological payoffs are relevant. With the Competition and No Competition results in hand, we next compare these results to our Monetary treatment and document how the results in this treatment differ from treatments with psychological payoffs. In particular, we show that behavior in the Monetary treatment is much closer to the pooling equilibrium compared with the two other treatments with psychological payoffs. We next document the performance of markets with psychological payoffs both with and without competition and the welfare implications of introducing competition among sellers.

We conduct two types of analysis. We first analyze the aggregate data and then, in Section 7, look at the individual subject data to unearth the possibility that different subjects might have played different strategies belonging to different equilibria identified above. This last exercise allows us to recover the distribution of strategies used by Buyers and Sellers in No Competition and Competition treatments and learn about equilibria selection in these environments.

4.1 Approach to Data Analysis

Given our experimental procedures (described in Section 3), the Buyers and the Sellers state their full strategies and beliefs at the beginning of each block for ten 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 will 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 (for example 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 report p-value associated with it. We use standard classification of significance levels: *** indicates significance at 1% level, ** indicates significance at 5% level, and * indicates significance at 10% level.

Most of the analysis presented below focuses on the last 5 blocks of each experimental session. This is done because it is common to observe subjects learning the game by playing it, which is why the data from the first iterations of the game tends to be noisier as subjects are trying to figure out their strategies. By the second half of the experiment subjects have experienced the game many times and may have possibly converged to their preferred strategies. However, we also present subjects' behavior in the first 5 blocks in several figures and tables to highlight changes in subject behavior due to learning and experience in the game.

4.2 Buyer and Seller Behavior in the Competition and No Competition Treatments

In this section we investigate Buyer and Seller behavior in the presence of psychological payoffs with and without competition. We do this first to test the predictions of our theory concerning the strategies and beliefs of our subjects and second to provide the behavioral results needed to understand the impact of competition on markets with psychological payoffs.

Our experimental design allows us to observe a wide variety of behavior concerning Buyer and Seller strategies and beliefs, and compare them with the predictions from our theoretical model. For example, we observe the stated beliefs and purchasing cutoffs of Buyers for both messages m_0 and m_1 in each of our treatments. 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, i.e., types $\omega \leq \omega'(m_i)$ will purchase the product after observing m_i, while the remaining types will not. Additionally, in the Competition treatment, a Buyers revealed her choice of Sellers 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, i.e., 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.

4.2.1 Do Subjects Play Pooling Equilibria?

We start with analyzing whether subjects' behavior in either treatment is consistent with a Pooling equilibrium as proposed in Prediction 1 of Section 2.7. Remember, in a Pooling equilibrium, Buyers treat messages as uninformative, and, therefore, hold the same beliefs and act the same way irrespective of the message they receive. For each treatment, Table 2 presents Buyers' beliefs and cutoffs conditional on receiving two messages m_0 and m_1.

Table 2 finds little support for Pooling equilibrium behavior in either Competition or No Competition treatments. For example, column 4 reveals that for the last five blocks in the Competition treatment m_0 messages lead Buyers to believe that there was a 22% chance that the good being sold was a high-quality good. In contrast, upon receiving m_1 messages Buyers believed that there was a 77% chance that the good being sold was of high-quality. Similar results hold for the No Competition treatment. In each treatment, statistical tests reject that Buyers interpreted messages m_0 and m_1 the same way, and that they acted similarly upon receiving these messages ($p < 0.001$ in all pairwise comparisons in both treatments in the first 5 and in the last 5 blocks). In fact, there is not much difference in actions and beliefs of Buyers between the first and the second half of the experiment with the exception of a rise in beliefs in the No-Competition treatment after observing message m_1.[19] In other words, Buyers did feel that the messages they received had information

[19] Statistical tests confirm this: in the No-Competition treatment $p = 0.661$ ($p = 0.035$) for beliefs after receiving m_0 (m_1) and $p = 0.680$ ($p = 0.372$) for actions after receiving m_0 (m_1); in the Competition treatment $p = 0.570$

content and responded to those messages accordingly. This is in contradiction to Prediction 1.

Table 2: Message Interpretation and Actions of Buyers

	No-Competition		Competition	
	first 5 blocks	last 5 blocks	first 5 blocks	last 5 blocks
Buyers' belief after observing m_0	0.24 (0.03)	0.26 (0.06)	0.24 (0.06)	0.22 (0.03)
Buyers' belief after observing m_1	0.71 (0.02)	0.76 (0.03)	0.75 (0.04)	0.77 (0.03)
Buyers' cutoff after observing m_0	0.31 (0.02)	0.30 (0.03)	0.34 (0.02)	0.31 (0.03)
Buyers' cutoff after observing m_1	0.55 (0.04)	0.59 (0.05)	0.56 (0.04)	0.62 (0.04)

Notes: Average beliefs and cutoffs of Buyers are reported with robust standard errors in the parenthesis, where observations are clustered at the session level.

We turn next to Sellers' behavior. Figure 4 presents the average frequencies of m_1 messages sent by Sellers of different psychological types in the last 5 blocks of the experiment. The communication strategy of Sellers in the first 5 blocks looks very similar and is presented in the Supplementary Appendix. A Pooling strategy for Sellers would suggest that the messages sent should be independent of the quality of good being sold. If instead, there is any correlation between the quality of the good and the messages sent by the Sellers, then Buyers could obviously use this information and choose a buying strategy that was a best response. Thus to maintain a Pooling equilibrium, all messages must be uninformative, and consequently, there should be no correlation between the quality of the good sold and the messages sent. The data presented in Figure 4 confirms that Sellers are not playing a Pooling equilibrium. For example, it is clear that different types of Sellers sent different messages conditional on the quality of the good they had. Overall, across all types and treatments, when a good was of high quality Sellers almost always revealed its true quality, while, when it was of low quality, they often lied. The extent to which they lied depended on their psychological type. Across all Sellers in any treatment, however, it was not the case that messages were not informative.

To substantiate this observation statistically, we ran random effects LOGIT regression for each treatment, where the dependent variable was an indicator variable that takes value 1 if message m_1 was sent and zero otherwise, and the independent variables were a constant and a dummy variable representing the quality of the good (coded as 1 for high quality of the good and 0 otherwise). We clustered observations by session as described in Section 4.1. If messages were uninformative, we would expect that the coefficient for the dummy variable to be insignificantly different from zero. Our regressions indicate that we must reject the hypothesis of a lack of correlation between the messages sent and the quality of the good, in favor of a positive relationship between messages and good quality in both treatments and in both parts of the experiment. The estimated coefficient on the product quality dummy is significantly different from zero in all these regressions with $p < 0.001$.

Finally, there should be no purchases made in a Pooling equilibrium. This is clearly not the case. In fact we find that the good was sold to the Buyer 43% of the time in the last 5 blocks of the No-Competition treatment, and 57% of the time in the last 5 blocks of the Competition treatment. This information is summarized in Figure 5. On the basis of these results we feel confident to reject the idea that subjects, in aggregate, adhered to a Pooling equilibrium in our experiment. We summarize the conclusion of the discussion above in Result 1 below.

($p = 0.352$) for beliefs after receiving m_0 (m_1) and $p = 0.359$ ($p = 0.083$) for actions after receiving m_0 (m_1).

Figure 4: Communication Strategy of Sellers (last 5 blocks)

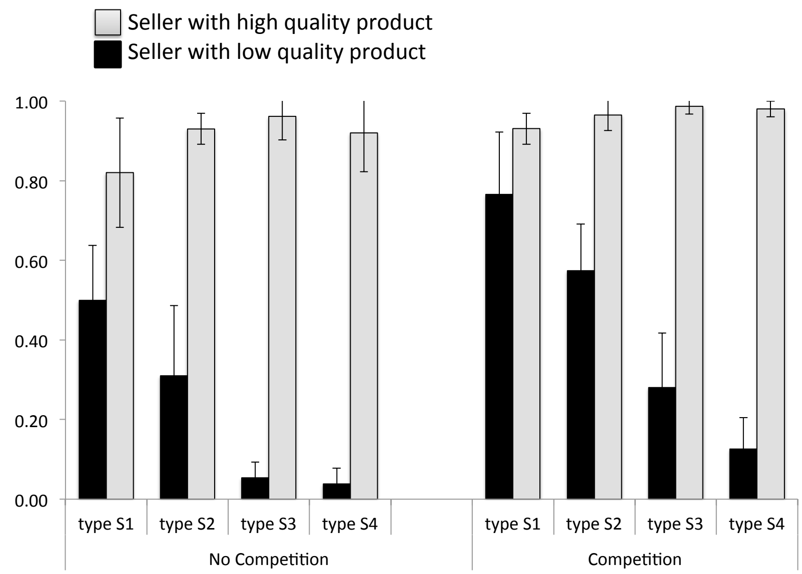

Notes: Average frequency of sending message m_1 is presented for each type of the Seller in each treatment in the second half of the experiment. 95% confidence intervals are computed using robust standard errors obtained by clustering observations by session.

Result 1: *In both the Competition and the No Competition treatment, Sellers strategically send messages, and Buyers believe that the messages provide information about the product quality and consequently use different purchasing strategies upon receiving different messages. Such a pattern indicates that Pooling equilibrium is not a good description of aggregate behavior in the two treatments with psychological payoffs.*

4.2.2 Do Subjects Play PIE1 or PIE2 Equilibria?

Having dispensed with the Pooling equilibrium we now turn our attention towards determining whether the observed behavior can support either of the partially informative equilibria we characterized earlier in Section 2.5. In particular, recall that a PIE1 and a PIE2 differ in the number Sellers with low quality product that lie in the communication stage.

We first focus on Buyers' beliefs and actions. Predictions 2, 3 and 3a specify the beliefs that Buyers should hold after receiving messages m_1 and m_0 and the corresponding cutoffs Buyers should set. Using the data summarized in Table 2, we reject the null hypothesis that average beliefs of Buyers after receiving message m_1 in the second half of the experiment in both treatments are not significantly different from 0.57 as suggested by PIE1 ($p < 0.001$). We cannot however reject the null that these beliefs are statistically different from 0.73 as suggested by the point predictions describing PIE2 ($p = 0.259$ in the No-Competition and $p = 0.236$ in the Competition treatments in the last 5 blocks). Thus, our results indicate that Buyers' beliefs upon receiving message m_1 seem to be consistent with PIE2 rather than PIE1 for both the treatments. At the same time, Buyers' cutoffs determining their purchasing decisions upon receiving message m_1 are lower than those predicted by PIE2 ($\omega^{\text{PIE2}}(m_1) = 1$) but higher than those predicted by PIE1 ($\omega^{\text{PIE1}}(m_1) = 0.51$). In the No-Competition treatment, using the data from the last 5 blocks, we reject the hypothesis that $\bar{\omega}(m_1) = 0.51$ and $\bar{\omega}(m_1) = 1$ ($p = 0.09$ and $p < 0.001$, respectively), where $\bar{\omega}(m_i)$ denotes the average observed cutoff for message m_i. In the Competition treatment using the data from the last 5 blocks, we also reject the hypothesis that $\bar{\omega}(m_1) = 0.51$ and $\bar{\omega}(m_1) = 1$ ($p = 0.01$ and $p < 0.001$).

Further, upon receiving message m_0 in either a PIE1 or a PIE2, Buyers should set their belief to zero and never buy the good, i.e., $b_1^B(m_0) = 0$, and $\omega(m_0) = 0$. Since these are corner solutions, all the deviations are going to be positive and will move the average Buyers' beliefs upon receiving m_0 and corresponding cutoff away from the prediction of zero. Thus, it is not surprising that we reject these predictions. Buyers' average beliefs in either treatment during the last 5 blocks were 0.26 and 0.22 for the No-Competition and Competition treatments, respectively, which were significantly different from zero ($p < 0.001$ in both treatments in the last 5 blocks).

The same is true about the purchasing cutoffs that Buyers set upon receiving an m_0 message. In the No-Competition treatment the mean of these cutoffs was 0.30, while it was 0.31 in the Competition treatment, both of which were significantly different from zero ($p < 0.001$). A perhaps more informative statistic regarding Buyers' beliefs and cutoffs conditional on receiving message m_0 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 5 blocks of the experiment, the majority of reported beliefs upon observing message m_0 are at most 5 percentage points away from zero. This happens in 63% of the time in the No-Competition treatment and in 77% of all cases in the Competition treatment. Similarly, in 61% of all cases Buyers' cutoffs are at most 5 percentage points away from zero in the last 5 blocks in the No-Competition treatment, while the same number is 74% in the Competition treatment.

There is a simple explanation 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. 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, i.e., be risk-averse, then we could observe Buyers holding the 'correct' beliefs but using lower cutoffs, which is precisely what our data indicates.

To test whether this explanation is supported by our data, 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 which paid one point for every point invested and a risky asset which 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 the full amount in the risky asset are either risk-neutral or risk-loving, while lower than full investment indicates that 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 indicates that there is 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 ($p = 0.068$).[20]

[20]To reach this conclusion we uses 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

We now ask how does a Buyer select a Seller in the Competition treatment? There are two distinct scenarios to consider: one in which both Sellers send the same message, and in the other, one Seller sends message m_0 while another one sends message m_1. In the first scenario, the Buyer has no basis to prefer one Seller over the other. Our data confirms that. In the last 5 blocks of the experiment, Buyers chose Seller 1 in such situation 53% of the times, but this fraction is not statistically different from 50% ($p = 0.527$). The same is true in the first 5 blocks of the experiment: upon observing identical messages, Buyers choose Seller 1 in 54% of the cases, which is not statistically different from 50% ($p = 0.233$). However, when Buyers got two different messages, in vast majority of cases they chose the Seller with message m_1 rather than m_0; this happened in 80% of the cases in the first 5 blocks and in 84% of the cases in the last 5 blocks of the experiment. Both of these fractions are significantly different from 50% ($p < 0.001$ in both cases).

Result 2: *In both the Competition and No Competition treatments, Buyers' beliefs and cutoffs upon receiving message m_1 (conditional on being risk-averse) are consistent with playing PIE2 where only type S1 Sellers with a low quality good lie in the communication stage. In addition, in the Competition treatment, Buyers choose the Seller who sends message m_1 if two messages are different, while they pick a Seller randomly if both messages are the same.*

Turning our attention to Sellers' communication strategies and Predictions 4 and 4a, we first examine the behavior of Sellers who own a high quality product. In both PIEs and in both treatments, high-quality Sellers should always tell the truth. Figure 4 shows that 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 product send message m_1 and look at whether our theoretical prediction (which says this likelihood should be one) 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 the types S2, S3, and S4 but not for S1.

The situation changes when we look at Sellers with low quality goods. Here the main difference in predicted behavior between PIE1 and PIE2 stems from the number of psychological types of Sellers that lie in the equilibrium. In PIE 1, both types S1 and S2 are expected to lie and send message m_1, while 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 and send message m_1, while other types (S2, S3, and S4) lie much less which is at least qualitatively consistent with PIE2. In 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), while types S3 and S4 lie much less. These results do not support either of our PIEs strictly, leading us to reject the hypothesis that either PIE was adhered to ($p < 0.001$ in all cases) in terms of Seller strategies. However, since we observe a monotonic decrease in the frequency of sending m_1 message by low quality Sellers as one moves from type S1 to type S4 in both treatments, Sellers in both treatments display behavior consistent with the weaker version of Prediction 4, namely, Prediction 4a. Statistical tests comparing the pairwise frequency of m_1 messages across types yield support for Prediction 4a.[21] What this analysis means is that the vast majority of Sellers with high-quality

consistent coefficients, correlations and standard errors and an estimator which is more efficient than standard instrumental variable techniques.

[21] Focusing on the Sellers with low quality products, in the last 5 blocks of the No-Competition treatment we

goods are truthful in their messages in both treatments, while Sellers with low quality goods are strategic in choosing their messages, taking into account the psychological costs of both lying and guilt sensitivity.

Result 3: *In both games with and without competition, Sellers with high quality goods tend to tell the truth in the communication stage of the experiment, while Sellers with low quality goods but higher values of guilt and lying sensitivity (types S3 and S4) lie less than those with lower values.*

4.2.3 Beliefs, Cutoffs, and Messages Between Treatments

Looking first at Buyers' beliefs and purchasing cutoffs, we see that they are remarkably similar both with and without competition in the last 5 blocks of the experiment (see Table 2). For example, during the last five blocks of the experiment after observing an m_0 message the mean belief of Buyers in the No-Competition treatment was 0.26 while it was 0.22 when there was competition. 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 0.30 and 0.31 for the No-Competition and Competition treatment, respectively ($p = 0.627$), while 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 contrast to Buyers' behavior, Sellers behave very differently when competition is introduced. Comparing behavior of Sellers between treatments, we observe a systematic difference between amount of lying in the presence and absence of competition. 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 4). For example, in the No-Competition treatment, while Seller with types S1, S2, S3, and S4 sent untruthful m_1 messages conditional on having a low quality good, 47%, 29%, 4%, and 2% of the time respectively, these percentages increased to 77%, 57%, 28%, and 13%, respectively, when there was competition. Pairwise comparisons between these fractions confirm the directional results evident in Figure 4 ($p = 0.015$ for types S1, $p = 0.006$ for types S2, $p = 0.005$ for types S3, and $p = 0.031$ for types S4).

An alternative way to support this claim is to use stacked regression analysis, where we analyze messages sent by the Sellers that own low quality products depending on their psychological type and the treatment they belong to (see Table 3). These regressions confirm the pattern documented above. First, Sellers with higher psychological disutilities from guilt and lying tend to lie less in the communication stage in both treatments. Second, there is a substantial increase in the amount of lies observed between the two treatments 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 4: *There is no statistical difference across our Competition and No-Competition treatments in terms of Buyers' beliefs and actions. At the same time, competition leads to a significant increase in the amount of lying that Sellers with low quality products engage in during the communication stage, irrespective of the guilt and lying sensitivity parameters of a Seller.*

---

observe significant difference between frequency of sending m_1 message by types S1 and S2 ($p = 0.006$), and between S2 and S3 types ($p < 0.001$), while there is no significant difference between frequency of sending m_1 for types S3 and S4. In the last 5 blocks of the Competition treatment, all pairwise comparisons between frequency of sending message m_1 of Sellers with low quality products with different psychological types are statistically significant with $p < 0.001$.

Table 3: Messages Sent by Sellers with Low Quality Products

	first 5 blocks	last 5 blocks
Seller's psychological type		
S2	-0.83*** (0.21)	-0.46*** (0.25)
S3	-1.61*** (0.26)	-1.29*** (0.41)
S4	-1.92*** (0.38)	-1.78*** (0.51)
Indicator for Competition treatment	1.08*** (0.19)	1.34*** (0.24)
Constant	0.009 (0.22)	-0.36 (0.28
# of observations	2178	2184
# of clusters	7	7
Log pseudo-likelihood	-1299.92	-1294.41

Notes: LOGIT regressions with dependent variable equals 1 if the Seller sent message m_1 and zero 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.

4.2.4 Higher-Order Beliefs

According to Prediction 5, in any equilibrium, we expect to see that the first-order beliefs of Buyers coincide with the second-order beliefs of Sellers and also coincide with the actual play of subjects for every possible message, which we will denote by $\bar{q}_H(m_i)$ i.e.,

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

In other words, in equilibrium all beliefs are confirmed. Note however, that while in either of our PIE's beliefs are expected to be confirmed, the beliefs consistent with each equilibria are different. In the analysis presented below we attempt to investigate whether the beliefs observed in our data are more consistent with one as opposed to another of our PIEs.

We start by comparing beliefs of Buyers and Sellers to those predicted in PIE1 and PIE2. Table 4 reports average beliefs of Buyers and Sellers observed in different treatments as well as the results of statistical tests comparing these beliefs with those predicted by PIE1 and PIE2. The predicted beliefs for PIE1 and PIE2 are listed at the top of the columns corresponding to PIE1 and PIE2, respectively.[22] As Prediction 2 suggests, we expect in our experiment the mean Buyer beliefs conditional on receiving an m_1 message to be 57% and 73%, respectively, for a PIE1 and PIE2 equilibrium and $b^1_B(m_0) = 0\%$ for each equilibrium when an m_0 is received.

The results in Table 4 are striking and are dramatically different for PIE1 and PIE2. In every single pairwise comparison, we reject the hypothesis that beliefs of either Sellers or Buyers correspond to the beliefs predicted by PIE1. In contrast, such a hypothesis can not be rejected for PIE2 when the m_1 messages were sent. For example, over the last five blocks of the experiment the PIE2 equilibrium predicts a mean first order belief of $\bar{b}^1_B(m_1)$=0.73 on the part of Buyers when they receive an m_1 message, and such beliefs were observed to be 0.76 and 0.77 for the No-Competition and Competition treatments respectively. As seen in Table 4 these beliefs are not significantly different from those predicted by PIE2.

[22] We have discussed Buyers' Beliefs in the previous section, in which we analyzed behavior of Buyers within each treatment. We reproduce some of this information in this table for completeness and ease of exposition.

PIE2 however fails in organizing our Buyer data when it comes to their responses after observing the m_0 messages. While theoretically $\bar{b}^1_B(m_1)$, and $\bar{b}^2_S(m_1)$ both should be 0, in our data they are strictly positive for both Buyers and Sellers. This divergence is puzzling: even though Sellers admit owning a low quality product, Buyers choose to ignore their warnings. This too will obviously have welfare consequences.

Table 4: Comparing Buyers' and Sellers' Beliefs with PIE1 and PIE2

	Data		PIE1: $b(m_0) = 0$, $b(m_1) = 0.57$		PIE2: $b(m_0) = 0$, $b(m_1) = 0.73$	
	$\bar{b}^1_B(m_i)$	$\bar{b}^2_S(m_i)$	$\bar{b}^1_B(m_i) = b(m_i)$	$\bar{b}^2_S(m_i) = b(m_i)$	$\bar{b}^1_B(m_i) = b(m_i)$	$\bar{b}^2_S(m_i) = b(m_i)$
first 5 blocks						
No-Competition						
message m_0	0.24 (0.03)	0.22 (0.03)	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
message m_1	0.71 (0.02)	0.73 (0.02)	$p < 0.001$	$p < 0.001$	$p = 0.272$	$p = 0.802$
Competition						
message m_0	0.24 (0.06)	0.25 (0.02)	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
message m_1	0.75 (0.04)	0.68 (0.02)	$p < 0.001$	$p < 0.001$	$p = 0.259$	$p = 0.889$
last 5 blocks						
No-Competition						
message m_0	0.26 (0.06)	0.20 (0.04)	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
message m_1	0.76 (0.03)	0.73 (0.02)	$p < 0.001$	$p < 0.001$	$p = 0.634$	$p = 0.063$
Competition						
message m_0	0.22 (0.03)	0.25 (0.02)	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
message m_1	0.77 (0.03)	0.70 (0.03)	$p < 0.001$	$p < 0.001$	$p = 0.236$	$p = 0.383$

Notes: This table reports average observed beliefs of Buyers and Sellers for each possible message in the first two columns, with robust standard errors reported in the parenthesis, where observations are clustered at the session level. The third and the fourth columns report the results of statistical tests comparing Buyers' and Seller's beliefs with those predicted by PIE1. The fifth and the sixth columns report the results of statistical tests comparing Buyers' and Seller's beliefs with those predicted by PIE2.

Table 5 adds one more piece of information to the beliefs story described above. In addition to listing the beliefs of Buyers and Sellers that were earlier presented in Table 4, Table 5 also depicts the actual probability that message m_i comes from the high quality Seller, $\bar{q}_H(m_i)$, and it presents statistical tests comparing Buyers' and Sellers' beliefs with each other as well as with the actual frequencies observed in the experimental sessions. Focusing on the last 5 blocks of our experimental sessions, we observe a few interesting patterns. First, casual observation of Buyers' and Sellers' beliefs suggest that these beliefs are quite similar to each other in both treatments conditional on the observed messages. For example, in the No-Competition treatment over the last 5 blocks, conditional on receiving an m_0 message the mean belief of Buyers about the probability of the good being high quality was 0.26 while the second order belief of the Seller about that belief was 0.20. For the m_1 message those beliefs were 0.76 and 0.73 respectively. The results for the Competition treatment were comparable with $\bar{b}^1_B(m_0) = 0.22$ and $b^2_S(m_0) = 0.20$ while $\bar{b}^1_B(m_1) = 0.77$ and $b^2_S(m_1) = 0.70$. Statistical tests confirm that there is no statistical difference between the beliefs conditional on message m_1 in No-Competition treatment ($p = 0.136$) and message m_0 in the Competition treatment ($p = 0.318$). There is a statistical difference, however, between Buyers' and Sellers' beliefs for m_0 in the No-Competition treatment and for m_1 message in the Competition treatment. Despite these statistical differences, however, the average beliefs of

Table 5: Buyers' and Sellers' Beliefs

	$\bar{b}^1_B(m_i)$	$\bar{b}^2_S(m_i)$	$\bar{q}_H(m_i)$	$\bar{b}^1_B(m_i) = \bar{b}^2_S(m_i)$	$\bar{b}^1_B(m_i) = \bar{q}_H(m_i)$	$\bar{b}^2_S(m_i) = \bar{q}_H(m_i)$
first 5 blocks						
No-Competition						
message m_0	0.24 (0.03)	0.22 (0.03)	0.07 (0.04)	$p = 0.218$	$p < 0.001$	$p < 0.001$
message m_1	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	0.24 (0.06)	0.25 (0.02)	0.22 (0.02)	$p = 0.680$	$p = 0.757$	$p = 0.052$
message m_1	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	0.26 (0.06)	0.20 (0.04)	0.08 (0.04)	$p = 0.009$	$p = 0.001$	$p = 0.002$
message m_1	0.76 (0.03)	0.73 (0.02)	0.71 (0.02)	$p = 0.136$	$p = 0.067$	$p = 0.247$
Competition						
message m_0	0.22 (0.03)	0.25 (0.02)	0.17 (0.02)	$p = 0.318$	$p = 0.167$	$p < 0.001$
message m_1	0.77 (0.03)	0.70 (0.03)	0.49 (0.03)	$p = 0.004$	$p < 0.001$	$p < 0.001$

Notes: This table reports average observed beliefs of Buyers and Sellers for each possible message in the first two columns. The third column reports the observed likelihood that message m_i comes from the high quality Seller. In all cells in the first three columns we report robust standard errors in the parenthesis, where we cluster observations at the session level. The last three columns report results of statistical tests comparing Buyers' and Seller's beliefs (fourth column), Buyers' beliefs and actual frequencies (fifth column), and Seller's beliefs and actual frequencies (sixth column).

the two groups are surprisingly close to each other. This is good news for Sellers, since in the belief task they were rewarded for how close their reported beliefs were from the beliefs stated by Buyers. Buyers, on the other hand, were rewarded for guessing the actual probability that message m_i comes from a Seller with high quality product, therefore, the relevant statistical test for these are reported in column 5. Except for message m_0 in the Competition treatment, our Buyers are doing quite poorly as they consistently overestimate the likelihood that message m_i comes from a Seller who owns a high quality product (compare beliefs in column 1 and column 3). The difference is especially pronounced in Competition treatment for message m_1 in which average Buyers' beliefs are almost 30 percentage points above actual ones. For comparison, in the No-Competition treatment, average Buyers' beliefs for m_1 are only 5 percentage points above the actual ones.

Result 5: *PIE2 outperforms PIE1 in organizing our data on first and second order beliefs of Buyers and Sellers. Specifically, we find Buyers' and Sellers' beliefs reported for message m_1 are statistically indistinguishable from those predicted by PIE2 in both treatments. We do find though that subjects report significantly higher beliefs for message m_0 compared to those predicted by PIE2. Sellers' second-order beliefs also match Buyers' first-order beliefs quite well, indicating that Sellers understand how Buyers interpret messages. However, Buyers remain somewhat clueless, consistently overestimating the meaning of both messages in both treatments. This effect is especially pronounced in Competition treatment for message m_1.*

5 Impact of Competition on Market Performance and Welfare in Treatments with Psychological Payoffs

With the results listed above we are in a position to understand how the behavior of our subjects across treatments affect market performance and wefare in our markets. Although in theory, competition does not introduce any new equilibria, there are some distinct ways in which competition can affect market performance. First in the Competition treatment, since there are two Sellers as opposed to one, there is a greater chance that at least one of them will be endowed with a high quality good. If Sellers use the same equilibrium communication strategies in both treatments, and Buyers interpret messages the same way in both treatments, then the Competition treatment should deliver higher payoffs for the Buyer than No-Competition treatment because the probability that at least one Seller will have a high quality good and send a strong message m_1 is higher than with one Seller. Of course this is an equilibrium result. Out of equilibrium, Sellers in the Competition treatment may feel compelled to lie more since it is a dominant strategy to send her m_1 message when receiving a high quality good and as a result those with low quality goods who send truthful m_0 messages will never be selected. Another reason why competition may lead to higher welfare is because competition may select a 'better' equilibrium (PIE2 rather than PIE1) from the perspective of the Buyer. However, the equilibrium selection story may work in an opposite direction as well: Buyers' welfare may decrease in the presence of competition if subjects settle on PIE1 or Pooling equilibrium in the Competition treatment as opposed to PIE2 equilibrium selected in the No-Competition treatment. It is then purely an empirical question as to which of these assertions are relevant for our experiment.

The question we ask in this section is how does competition affect the transaction rates and ultimately welfare in our markets? Figure 5 depicts the frequency of Buyers' purchasing decisions in each treatment as well as the quality of the purchased good.

As Figure 5 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 as it captures behavior of subjects after they have had the time to learn the game and converge to a stable strategy. Figure 5 shows two main results. First, competition increases trade: over the last 5 blocks of the No-Competition treatment only 43% of the products available were actually sold, while this percentage increases to 57% when there is a competition 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 5 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, $p = 0.098$).[23]

Table 6 reports regressions that investigate the effect of competition on Buyers' and Sellers' welfares. To make comparison between Sellers' payoffs in the two treatments valid, we focus on the payoffs of the selected Seller in the Competition treatment, since the non-selected Seller earns a fixed payoff of zero. Also, in these regressions we abstract away from the payoffs subjects accumulated in the belief-guessing task and focus only on the tree game payoffs.

A few interesting patterns emerge from this analysis. Although there are no differences in the average payoffs of Buyers in the first half of the experiment, Buyers earn significantly less in the Competition than in the No-Competition treatment by the end of the experiment (last 5 blocks). In

[23] These quantities are not statistically significant in the first 5 blocks: $p = 0.223$ and $p = 0.840$ in two comparisons, respectively.

Figure 5: Aggregate Outcomes, by Treatment

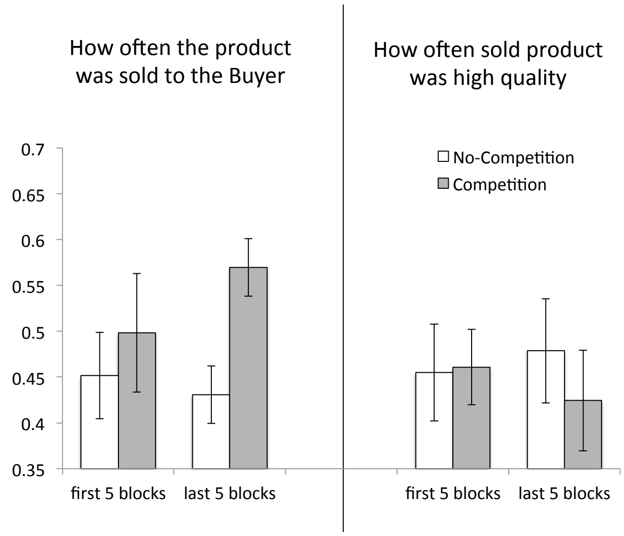

Notes: The left panel depicts purchasing frequency by treatment. The right panel depicts the likelihood that the product was high quality conditional on the product being purchased. Bars indicate 95% confidence intervals using robust standard errors, which are computed by clustering observations by session.

Table 6: 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 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 subjects accumulated for guessing beliefs tasks. Standard errors are clustered at the session level. ** indicates significance at 5% level, while *** indicates significance at 1% level.

other words, Buyers suffer from the presence of competition. This result contradicts Prediction 6, which states that in any partially informative equilibrium, the payoffs of the Buyers are higher when there is a competition between Sellers. The Sellers experience a similar effect which is pronounced right from the start of the experiment: both in the first half and in the last half, selected Sellers in the Competition treatment earn less than Sellers in the No-Competition treatment. These effects are statistically significant and meaningful in terms of eventual payoffs our experimental subjects earned in the experiment.

Which types of Sellers and Buyers suffer the most from the presence of competition? Table 7 presents average payoffs of Sellers and Buyers broken down by their types in each of the treatment. We focus here on the last 5 blocks of the experiment and refer the reader to the Supplementary Appendix for the same statistics regarding the first 5 blocks of the experiment.

Table 7: 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 for each type in the last 5 blocks of the experiment and the robust standard error in the parenthesis, where standard errors are clustered at the session level. The last column reports the result of a statistical test comparing payoffs for a fixed type of Buyers or Sellers in the two treatments. The test is performed using Random Effects GLS regressions, in which we regress the payoffs of interest on a constant and a dummy variable indicating one of the treatments, while clustering observations by session. Statistical difference between treatments is assessed by looking at the significance of the estimated dummy coefficient. We use standard convention of * indicating significance at 10% level, ** indicating significance at 5% level, and *** indicating significance at 1% level and report p-value associated with estimated dummy variable in the parenthesis.

As Table 7 shows Sellers who own the high quality product earn the same average payoffs in No-Competition and in Competition treatments irrespective of their psychological type. It is the Sellers with low quality goods who suffer from the competition, as 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 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 which are not statistically different from zero, which is precisely what the non-chosen Sellers get. As for the Buyers, it is those with higher disappointment aversion who suffer the most losses from competition between Sellers.

Results 6: *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.*

5.1 The Mechanism Through Which Competition Affects Market Outcomes

Our results in the Competition and No-Competition treatments paint a clear picture of the effects of competition among Sellers. Buyers unequivocally suffer in the presence of competition, as they are more likely to purchase lower quality products. Consequently, they earn lower average payoffs (Result 6).

The decreased welfare in the Competition treatment is the result of a number of factors. First, while Buyers tend to use the same strategies in both treatments, Sellers with low quality products lie more often in the Competition treatment (Result 4). Moreover, Buyers on average hold correct beliefs about the communication strategies of the Sellers in the No-Competition treatment, but consistently overestimate how likely message m_1 comes from the high quality Sellers in the Competition treatment (Result 5). In other words, Buyers fail to discount the meaning of m_1 messages in the Competition treatment relative to the No-Competition treatment, although they should have, given that Sellers with low quality products tend to lie more there.

Figure 6: Evolution of Buyers' Beliefs and Quality of the Product for Message m_1

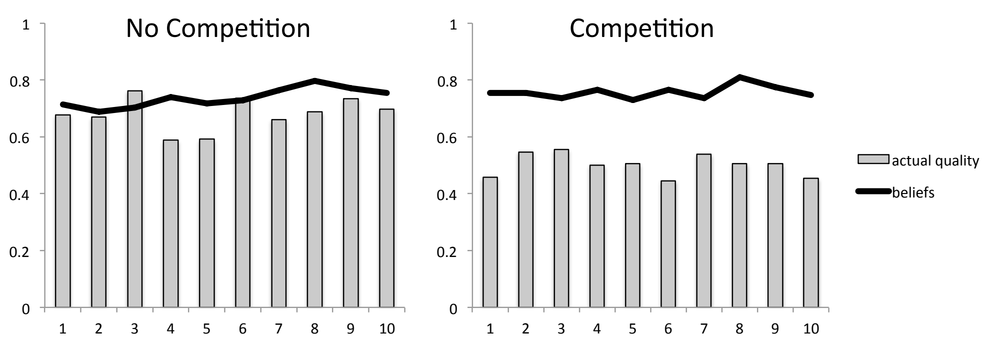

Notes: The lines depict the average Buyers' beliefs in each block in each treatment, while the bars depict the average frequency that message m_1 comes from the high quality product. The data is broken by the treatments with No-Competition in the left panel, and Competition in the right panel.

To drive home this point, Figure 6 presents the evolution of the average Buyers' beliefs (depicted by a line) and actual frequency that message m_1 comes from Seller with high quality products (depicted by a bar), block by block in each treatment. In the No-Competition treatment, our Buyers are holding beliefs which are very close to the actual meaning of the m_1 message given Sellers' strategies. However, in the Competition treatment, there is a big gap between what Buyers think message m_1 means and what it actually means. In other words, in the Competition treatment Sellers are successfully fooling Buyers by sending message m_1 more often when they actually have low quality products and Buyers are not adjusting their beliefs accordingly.

6 The Monetary Treatment: Games with No Induced Psychological Payoffs

Up until now we have only compared the performance of our markets with the predictions of theory. However, our Monetary treatment allows us to also compare these markets empirically since in the Monetary treatment no psychological payoffs are imposed on our subjects. Hence, by comparing the behavior of subjects across the Monetary and No-Competition treatments we can asses the impact of imposing psychological payoffs.

As we discussed in Section 2.1, the game with only monetary payoffs (see Figure 1) has a unique equilibrium in which trade does not occur because Buyers cannot distinguish the high quality products from the low quality ones. This is a pooling equilibrium, in which Sellers' messages are not informative, and, thus, must be ignored by the Buyers. The introduction of psychological forces in this market may boost trade, however, if subjects reject the pooling equilibrium in favor of one of the two PIE's that exist when psychological payoffs are introduced. What we show in our analysis below is that behavior in the Monetary treatment, while not strictly consistent with all aspects of a pooling equilibrium, is significantly closer to it than the behavior we witnessed in the No Competition treatment.

The failure of behavior to exactly mimic that expected in the pooling equilibrium is understandable. First, while a pooling equilibrium makes some distinct predictions like no trade, its predictions about message strategies for the Sellers and beliefs for the Buyers are undefined. For example, in a pooling equilibrium sellers might send all m_1 messages, or all m_0 messages, or each message half of the time, etc. All that matters is that the quality of the good is not correlated with the message in a strong enough manner so as to be exploited by the Buyers in a profitable way. Further, beliefs for the Buyer may also not be unique since they are free to hold a variety of beliefs as long as these beliefs do not lead them to think that it is profitable to buy the good when receiving either message. One way to do this is to leave the beliefs the Buyers hold unchanged from their prior no matter what message they observe. Finally, on a behavioral level, it is not clear what type of behavior makes sense to subjects when they suspect they are being lied to. As one subject told us, "I bought when they sent the m_0 message and did not when they sent the m_1 message because I thought they were lying to me all the time so when they told me the good was of high quality I did not buy but when they said it was of low quality, I did. I did the opposite of what they said because they always lied". While this strategy has severe logical limitations, it does signal that understanding that one is being lied to and hence should not trust the Seller does not imply a unique buying strategy - at least to some subjects.

Finally, we might observe truth telling behavior on the part of our Sellers in the Monetary treatment (and we do) simply because while we did not induce any penalties for lying on them, they might have brought these norms with them into the lab. The fact that we observe significantly different behavior in our Competition and No Competition treatments and in the direction of our theoretical predictions suggests, however, that we were successful in inducing lying and guilt aversion in these treatments.

To examine behavior in the Monetary treatment compared with No Competition treatment, consider the individual strategies of Sellers and Buyers separately. What we are interested in here is whether subjects use strategies that are more consistent with pooling equilibrium strategies in the Monetary as opposed to the No Competition treatment. By a pooling equilibrium strategy for Sellers we mean choosing the same message irrespective of the quality of the product 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 irrespective of the received message in the Monetary treatment, and choosing the two 'similar' purchasing cutoffs conditional on either message, where similar means that cutoffs are no more than 0.05 away from each other. Note for the Monetary treatment this definition is very conservative for Sellers since it only considers behavior consistent with a pooling equilibrium if the Seller sends the same message no matter what the quality of the good received is while we know that there are a multitude of other strategies that might be consistent with pooling equilibrium behavior, i.e., any behavior which is such that the messages sent are not sufficiently informative.

Table 8 presents the frequency with which our sellers and buyers used the pooling equilibrium (as defined by our conservative definition) in the Monetary and No Competition treatments.

Table 8: Individual Freq of Pooling Eq Strategies in Monetary and No Competition treatments

	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%	4%	23%	25%
Monetary vs No Competition	$p < 0.001$	$p < 0.001$	$p = 0.008$	$p < 0.001$

Notes: Playing pooling equilibrium for Sellers entails choosing the same message irrespectively of the quality of the product in the Monetary treatment, and irrespectively of both the product quality and one's psychological type in the No Competition treatment. Playing pooling equilibrium for the Buyers entails making the same purchasing decision irrespectively of the received message in the Monetary treatment, and choosing the two 'similar' purchasing cutoffs, where similar means that cutoffs are no more than 0.05 away from each other.If we increase the band to 0.1, then there are 27% of Buyers who play pooling equilibrium strategy in the first half of the No Competition treatment, while there are 28% in the second half.

As is apparent from Table 8, in the second half of the experiment, 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. That is, more than 60% of Sellers send the same message irrespective of the quality of the product they own, and, 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.

In addition to the strategies used by our Buyers and Sellers we find support for differences in the beliefs held by the Buyers upon receiving the m_1 message across the two treatments. For example, in the second half of the Monetary treatment, after receiving message m_1 the average Buyers' belief that this message is sent by the Seller with a high quality product is 0.43 which is significantly lower than 0.76 belief observed in the No Competition treatment ($p < 0.001$).[24] Note that with these beliefs it is rational to buy the good when receiving the m_1 message in the No Competition treatment and to not do so in the Monetary treatment. In other words, if subjects were fully rational and held these average beliefs, then we should observe a no trade equilibrium in the Monetary but not in the No Competition treatment. However, heterogeneity exists in that

[24]The average Buyers' belief after observing message m_0 in the second half of the Monetary treatment is 0.10, while the corresponding average belief in the No Competition treatment is slightly higher and equal to 0.26.

Table 9: Correlation b/w High Quality of the Product and Message $m1$ in Monetary and No Competition treatments

	First 5 Blocks	Last 5 Blocks
Monetary treatment	0.39^{***}	0.32^{***}
No Competition treatment	0.61^{***}	0.65^{***}

Notes: To compute this correlation we look at strategies specified by Sellers for each block of the experiment. *** indicates significance at 1% level.

even in the Monetary treatment, there are some Buyers who believe that messages are informative and take them at face value, i.e., purchase the product when they receive message m_1 and don't purchase when they receive message m_0.[25] This fraction is 39% in the second half of the experiment, and this is why we observe positive trade frequencies in the Monetary treatment.

The results just discussed dealt with beliefs. One could ask in addition whether buying in the No Competition and not buying on the Monetary treatment conditional on receiving an m_1 message is a best response to the behavior of the Sellers. In other words, were the Sellers significantly more honest in the No Competition treatment as compared to the Monetary treatment as to warrant buying conditional on receiving an m_1 message in one but not in the other? We will answer this question in two ways. First we will look at the correlation between product quality and m_1 messages and see if they are such that they can be relied on in the No Competition treatment and not the other. Second we will look at the actual conditional probability of a high quality good being sold conditional on an m_1 message be sent in the population of Sellers to see if a Buyer's posterior probability suggests buying in the No Competition treatment and not doing so in the Monetary treatment.

To start with correlations, consider Table 9 which presents the correlation between messages and product quality in both treatments. While the correlations in both treatments are significantly different from 0, they are far greater in the No Competition treatment indicating that they are far more informative there than in the Monetary treatment. In addition, note that while there is learning in the Monetary treatment in that the correlation decreases across the first and last blocks, the correlation actually increases in the No Competition treatment suggesting that the messages received are more informative as time goes on.

Finally one can ask whether, given the different behavior of Sellers across these two treatments is it a best response for Buyers to buy conditional on receiving an m_1 message in the No Competition treatment but not so in the Monetary treatment? To answer this question we simply look to the data and calculate the observed probability of the good being of high quality given the receipt of an m_1 message in both treatments and 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 5 so buying conditional on receiving an m_1 message would have to yield more than 5 in order to be profitable. Since $Pr[\text{low quality}|m_1] = 0.519$ and 0.493 during the first and last five blocks in the Monetary treatment, we see that buying conditional on receiving the m_1 message while marginally profitable

[25]The existence of Buyers who blindly believe messages of the Sellers is consistent with Buyers holding beliefs that Sellers (and other people in this world) have psychological motives that shape their behavior, i.e., inner guilt and lying aversion that prevents Sellers from lying even when doing so increases their material payoffs.

in the first five blocks is marginally unprofitable after subjects learn in the last five blocks:

$$\mathbb{E}\Pi^{\text{Buyer}}_{\text{first 5 blocks}}|m_1 = 10 \cdot Pr[\text{high quality}|m_1] + 0 \cdot Pr[\text{low quality}|m_1] = 10 \cdot 0.519 = 5.19 > 5$$

$$\mathbb{E}\Pi^{\text{Buyer}}_{\text{last 5 blocks}}|m_1 = 10 \cdot Pr[\text{high quality}|m_1] + 0 \cdot Pr[\text{low quality}|m_1] = 10 \cdot 0.493 = 4.93 < 5$$

For the No Competition treatment, however, just the opposite is the case since the conditional probability of a good being high quality given an m_1 message is 0.656 and 0.705 for their first and last five blocks of the No Competition treatment and hence,

$$\mathbb{E}\Pi^{\text{Buyer}}_{\text{first 5 blocks}}|m_1 = 10 \cdot Pr[\text{high quality}|m_1] + 0 \cdot Pr[\text{low quality}|m_1] = 10 \cdot 0.656 = 6.56 > 5$$

$$\mathbb{E}\Pi^{\text{Buyer}}_{\text{last 5 blocks}}|m_1 = 10 \cdot Pr[\text{high quality}|m_1] + 0 \cdot Pr[\text{low quality}|m_1] = 10 \cdot 0.705 = 7.05 > 5$$

In other words, while subjects in the Monetary treatment learn to be skeptical of the behavior of Sellers, just the opposite seems to be the case for the Buyers in the No Competition treatment since they appear to trust the Sellers more and more as time goes on and actually should given their behavior since their messages become more reliable.

Finally, we can ask whether the different behavior of Buyers and Sellers across these treatments led to differences in the frequency of trade. Table 10 depicts trade frequencies in the Monetary and No Competition treatments. While 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 than in the No Competition treatment.

Table 10: 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	$p = 0.813$	$p = 0.003$	

Notes: We report the average trade frequencies in each treatment and statistical tests assessing whether these frequencies are different in the first half and the second half of each treatment (p-values reported in the last column) and whether these frequencies are different in the two treatments conditional on the first/second half of the experiment (p-values reported in the last row).

Our results here emphasize one salient feature of our Monetary treatment which is that not only does learning exist in that treatment but it moves behavior in the direction of the theoretical prediction, i.e., the pooling equilibrium. The fact that no learning occurred in the No Competition treatment or that when it did it move in the direction away from pooling equilibrium suggests that the lessons of the Monetary treatment might have been easier for our subjects to distill. Finally, the fact that some of our results in the Monetary treatment deviated from that expected in a perfect pooling equilibrium points out that while we were able to make lying and guilt more salient in our No Competition and Competition treatments, imposing no penalties for lying or guilt does not mean that those emotions do not exist. Some subjects simply do not like to lie and feel guilty if they do so that a model that assumes that people are capable of lying to or misleading others without feeling guilty (as conventional models do) may not be realistic.

Result 7: *Market outcomes and behavior of Sellers and Buyers in the Monetary treatment is 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 are then ignored by the majority of Buyers. This results in trade frequencies which are significantly lower in the Monetary than in the No Competition treatment, in which Sellers tend to use informative communication strategies and Buyers respond to messages. In other words, the introduction of psychological payoffs positively affects trade in the markets and shifts behavior towards some information transmission.*

7 Individual-level Analysis of Buyers' and Sellers' Strategies in the No Competition and Competition Treatments

In this section we look at individual behavior in an attempt to recover the distribution of strategies used by Buyers and Sellers in the No Competition and Competition treatments. This exercise is informative as it speaks to the equilibrium selection issue we brought up earlier.

A Seller's strategy consists of specifying eight messages: one for each of the eight psychological type-product quality pairs. In Table 11 we present a breakdown of our Sellers' strategies in each treatment. We treat strategies reported by Sellers in each block as an independent observation. This allows us to capture learning behavior across blocks, as subjects might change their strategies based on their experiences from previously played blocks.

Table 11: Sellers' Strategies in the Communication Stage

		first 5 blocks		last 5 blocks	
		No-Comp	Comp	No-Comp	Comp
	Total # of obs	100%	100%	100%	100%
	Sellers with $q = q_H$ and all psycho types send m_1	77% of which	89% of which	78% of which	90% of which
TRUTH	and $q = q_L$ with all psycho types send m_0	46%	20%	50%	21%
PIE2	and $q = q_L$ with S1 send m_1, others send m_0	17%	20%	17%	13%
PIE1	and $q = q_L$ with S1, S2 send m_1, others send m_0	26%	38%	25%	39%
	and $q = q_L$ with S1, S2, S3 send m_1, S4 sends m_0	4%	7%	1%	13%
POOL	and $q = q_L$ with with all psycho types send m_1	5%	8%	3%	9%
	and remaining observations	2%	6%	3%	5%

Notes: In this table we treat a strategy of a Seller in a block as an independent observation.

In all three equilibria, Sellers who own high quality products are expected to send truthful message m_1. Further, depending on the behavior of Sellers with low quality products, we can attribute these strategies to one of the considered equilibria, or to non-equilibrium behavior. Table 11 shows that most of the observations can be classified into three types of behavior: TRUTH (Sellers revealing the quality of their product truthfully, which is not part of any equilibrium strategy), PIE1 or PIE2. These three strategies emerge as the most commonly played strategies right from the start of the experiment and remain so till the end of the experiment. However, we find important difference in behavior between the two treatments especially after subjects had gained some experience with the game. In the last 5 blocks of the No-Competition treatment, the most common strategy used by the Sellers is TRUTH, i.e. the one in which messages truthfully reflect the qualities of the products owned by the Sellers. Such a strategy is observed in 50% of the cases in which Sellers with high quality products send message m_1. The fraction of Sellers telling the truth is significantly lower in the Competition treatment (only 21%). The most commonly used strategy in the Competition treatment is the PIE1 equilibrium strategy, in which Sellers with low quality products and psychological types of S1 and S2 lie and send an m_1 message, while the remaining types of Sellers with low quality products send an m_0 message. There is also a significant fraction of Sellers in both treatments who play the PIE2 equilibrium (17% in the No-Competition and 13% in the Competition treatment) in the last 5 blocks of the experiment. These results are consistent with the aggregate behavior of Sellers analyzed above, i.e., Sellers with low quality products lie much more in the Competition treatment than in the No-Competition treatment.

To classify Buyers' strategies, we use Buyers' cutoffs reported in each block of the experiment, and instead of the point predictions we use the qualitative features of different equilibria described in Section 2.5. We start by classifying Buyers' strategies into those that play Pooling equilibrium and those that play partially informative equilibria (the first two rows in Table 12). Buyers who set very similar cutoffs for both m_1 and m_0 messages, i.e., cutoffs that are less than 10 percentage points apart, are characterized as playing a Pooling strategy since they essentially behave the same way irrespective of the received message. On the contrary, Buyers who set cutoff for an m_1 message at least 10 percentage points higher than cutoff for an m_0 message are classified as playing a PIE. Distinguishing which PIE a Buyer is playing is a more complicated task, since as we argued in Section 4.2.1, risk attitude might affect purchasing cutoff for an m_1 message. However, if one adheres to the assumption of risk-neutrality, then we can say that a Buyer plays PIE1 if she sets the cutoff for message m_1 close to 0.51 (between 0.4 and 0.6 to allow for some small noise), while a Buyer plays PIE2 if she sets the cutoff for message m_1 close to 1 (above 0.8 to allow for small noise.[26]

Our data on Buyer behavior reveals a few interesting patterns. First, Table 12 shows that by the end of the experiment, only about a quarter of Buyers in each treatment play a Pooling strategy. In fact, Buyers in the Competition treatment play this strategy less and less as they experience the game: the fraction of those who play a Pooling strategy decreases from 40% in the first 5 blocks to 26% in the last 5 blocks. Second, in both treatments, we observe quite a lot of heterogeneity in terms of the cutoff that Buyers set after observing an m_1 message. This might be driven by differences in the risk attitudes of our experimental Buyers. Despite this heterogeneity, the vast majority of these cutoffs are quite high (above 0.6) which is consistent with playing the most informative PIE, i.e., PIE2, and being a risk averse. Third, if one insists on risk neutrality, then Buyers rarely play PIE1 strategy in either of the treatments (less than 15% in both treatments), while they play PIE2 strategy more often in the Competition than in the No-Competition treatment ($p = 0.057$).

In addition, Buyers' beliefs indicate that they feel that messages are more informative in the Competition than in the No-Competition treatment. This explains why m_1 cutoffs are generally higher in the Competition than in the No-Competition treatment. Recall that we have discussed the informativeness of different types of equilibria in Observation 1 (Section 2.5). According to this

[26] We focus on the purchasing cutoffs that Buyers report for an m_1 message, since this is what distinguishes different types of partially informative equilibria. All PIEs predict that purchasing cutoffs, upon observing an m_0 message should be zero. Despite that, most of our Buyers chose strictly positive purchasing cutoffs upon observing an m_0 message. This is consistent with the fact that some Sellers with high quality goods chose to send an m_0 message.

Table 12: Buyers' Purchasing Strategies

		first 5 blocks		last 5 blocks	
		No-Comp	Comp	No-Comp	Comp
	Total # of obs	100%	100%	100%	100%
POOL	$\lvert\bar{\omega}(m_1) - \bar{\omega}(m_0)\rvert < 0.1$	27%	40%	28%	26%
PIE	$\bar{\omega}(m_1) \geq \bar{\omega}(m_0) + 0.1$	57%	50%	55%	63%
		of which	of which	of which	of which
	$\bar{\omega}(m_1) < 0.4$	9%	0%	3%	1%
PIE1	$0.4 \leq \bar{\omega}(m_1) < 0.6$	15%	5%	11%	1%
	$0.6 \leq \bar{\omega}(m_1) < 0.8$	36%	60%	39%	41%
PIE2	$0.8 \leq \bar{\omega}(m_1)$	38%	35%	47%	57%

Notes: In this table we treat a strategy of a Buyer in a block as an independent observation.

observation, the informativeness is defined as the difference in beliefs upon observing an m_1 and an m_0 messages, and the most informative equilibrium is PIE2, while the least informative one is the Pooling equilibrium. Following this observation, in Figure 7 we plot the cumulative distribution of observed informativeness of messages in each treatment.

Figure 7: Observed Informativeness of Messages, by treatment

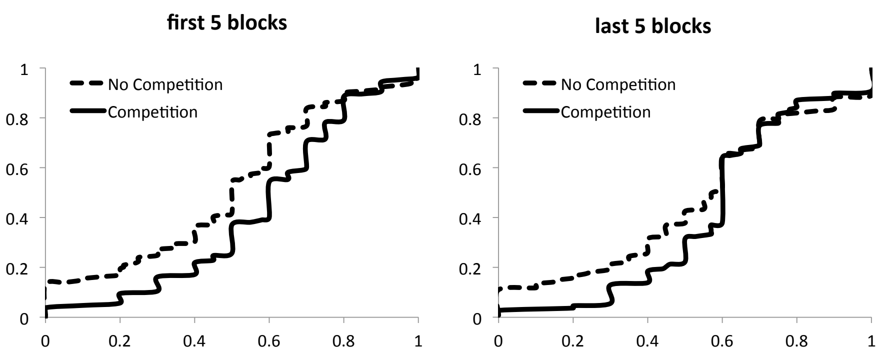

Notes: Observed informativeness measures the difference between a Buyer's belief upon observing an m_1 message and an m_0 message, where a belief upon observing message m_i indicates the likelihood that this message m_i comes from a Seller with a high quality product. We treat a pair of beliefs for a Buyer in a given block as an independent observation.

As Figure 7 shows, our Buyers felt that competition between Sellers increases the informativeness of messages. This is true from the start of the experiment. The cumulative distribution of informativeness in the Competition treatment is shifted to the right of the distribution in the No-Competition treatment for the vast majority of our observations in both the first 5 and the last 5 blocks of the experiment. The exception being Buyers who believe that Sellers always tell the truth in their messages (these Buyers report the difference in beliefs being close to 1). These

Buyers are mistaken however, since both at the individual as well as at the aggregate level, Sellers lie significantly more in the Competition than in the No-Competition treatment, and, even though in the Competition treatment Buyers are able to select one of the Sellers to buy from, the difference in the informativeness of messages in this treatment is significantly smaller than the one in the No-Competition treatment (see Table 5).

8 Conclusions

In this paper we studied the impact of competition on the welfare of buyers and sellers in a market where agents have psychological payoffs. More precisely we looked at sellers who suffer a cost both when they lie and/or mislead buyers into buying subpar goods while buyers suffer from disappointment whenever they are tricked into buying the subpar good.

In contrast to previous work on psychological games we induced the costs of lies, guilt and disappointment into the payoff functions of our subjects. Doing so allows us to control these psychological payoffs experimentally and evaluate their comparative static effects. As we have shown, our experimental design was successful in manipulating the payoffs faced by our subjects and hence achieved our goal. Of course one might suggest that we maybe inducing psychological payoffs where none might otherwise exist, i.e., we make subjects guilt averse or lying averse when they might not be so in real life. This does not do an injustice to our results, however, because either our subjects are more averse to lying and guilt than we make them, in which case our results furnish a lower bound on the impact of such emotions, or they are less averse, in which case we have controlled their emotions precisely for the game at stake. Importantly, the aim of our experiment is not to measure people's lying costs or guilt aversion but to see the impact of those factors on the behavior of our subjects and how they can help or harm in achieving equilibrium different than the typical pooling equilibrium.

Our results indicate two broad points about the workings of competitive markets in the presence of psychological payoffs. First, the pressure of competition, especially in a winner-take-all situation, encourages sellers to misrepresent and lie more, even if they suffer from psychological payoffs. This is not true in a situation where sellers do not compete for a buyer with other sellers. This seller propensity to lie more, however, is encouraged by the behavior of our buyers. For example, buyers seem to genuinely believe that competition is necessarily beneficial for their welfare, and consequently trust seller messages more often than in a situation without competition. Strangely enough, even with feedback and experience the buyer naivety does not go away. Consequently, sellers take advantage of such blind faith on the buyers' part and peddle lower quality products indiscriminately. Furthermore, although sellers with high psychological costs (such as our types S3 and S4) suffer more compared to the sellers in the no competition game with the same types, they are no worse off on an average compared to an unmatched sellers.

As mentioned in our introduction, the implications of our experiment extend beyond simple markets. In particular, it is not uncommon to see politicians lie through their teeth, and keep doing so repeatedly. Citizen's who exaggerate the ability of political competition to weed out such lies and bad agents might be at the heart of such continued behavior.

References

[1] 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.

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

[3] 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.

[4] 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.

[5] Cason, Timothy, Shremeta, Roman, and Jingjing Zhang (2012). "Communication and efficiency in competitive coordination games," *Games and Economic Behavior* 76(1): 26-43.

[6] Charness, Gary, and Martin Dufwenberg (2006)."Promises and Partnership," *Econometrica* 74: 1579-1601.

[7] Charness, Gary, and Martin Dufwenberg (2011). "Participation," *American Economic Review* 101: 1213-39.

[8] Crawford, Vincent, and Joel Sobel (1982). "Strategic Information Transmission," *Econometrica*, 50(6): 1431-1451.

[9] 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.

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

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

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

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

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

[15] Gillen, Ben, Snowberg, Erik, and Leeat Yariv (2018). "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study," *Journal of Political Economy*, forthcoming.

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

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

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

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

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

[21] Keck S, Karelaia N (2012). “Does competition foster trust?,” *Experimental Economics*, 15: 204-228.

[22] Jin, Ginger, Luca, Michael, and Daniel Martin (2017). “Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure,” Working Paper.

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

[24] Oprea, Ryan, Gary Charness, and Dan Friedman (2015). “Continuous Time and Communication in a Public-goods Experiment,” *Journal of Economic Behavior and Organization*, forthcoming.

[25] 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.

[26] 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.

[27] Vanberg, Christoph (2008). “Why do people keep their promises? an experimental test of two explanations,” *Econometrica*, 76(6): 1467-1480.