ESSAYS IN STRATEGIC INFORMATION TRANSMISSION

VALERIA BURDEA



Thesis submitted to The University of Nottingham
for the Degree of Doctor of Philosophy

School of Economics
University of Nottingham

September 2018

Valeria Burdea: *Essays in Strategic Information Transmission,* PhD in Economics, Supervised by Dr Maria Montero and Professor Martin Sefton © September 2018

As you are reading these words, you are taking part
in one of the wonders of the natural world.
For you and I belong to a species with a remarkable ability:
we can shape events in each other's brains with exquisite precision.

— Steven Pinker, *The Language Instinct* (1994)

ABSTRACT

This thesis contributes to the literature on strategic information transmission through providing new insights into strategic behavior under several communication frameworks. The thesis is comprised of five chapters.

Chapter 1 provides an overview of the main topics and research methodologies presented in the next three chapters.

Chapter 2 reports experiments in sender-receiver games with partially verifiable messages. We explore the role of verification control on strategic behavior and final outcomes. We find that behavior is closer to theoretical predictions when senders have verification control. However, the setting in which receivers have control over the verification action promotes higher average payoffs despite behavior being noisier.

Chapter 3 investigates the effect of cheap talk on the interpretation of partially informative verifiable communication (evidence). Using experimental data from sender-receiver games similar to those in chapter 2 we find that cheap talk distorts receivers' appraisal of evidence. Specifically, it makes receivers depart from theoretical predictions more than they do in a treatment where cheap talk is not present.

Chapter 4 explores the effect of cheap talk in a sender-receiver game with sender state independent preferences and a state space with non-uniform distribution. This type of state space leads to the interests of the players to be either more likely aligned, or more likely conflicting. Using a simple theoretical analysis we show that if receivers dislike taking an action that is profitable for a deceiving sender ("sucker aversion"), communication can harm when interests are more likely aligned. However, when interests are more likely conflicting, communication helps due to senders' "lying aversion". We run experiments to test these predictions and find no effect of cheap talk when interests are more likely aligned. When interests are more likely conflicting, communication has a positive effect on payoffs but not due to the hypothesized mechanism.

Chapter 5 provides a summary of the previous chapters' results as well as a discussion which points out the limitations of the studies. Finally it delineates welcomed directions for future research.

How do you go about searching for
knowledge you are not aware of?
If you should meet it, how will you know
that this is the thing you did not know?

— Plato, *Meno* (80d)

ACKNOWLEDGMENTS

It is fair to say that throughout my PhD studies, I rarely felt that I was purposefully working towards their completion. Rather, I have had the fortune to naturally stumble upon bits and pieces that I found interesting and that were deemed relevant for enlarging our (economic) knowledge. Perhaps the most important lesson I have learned during this time is that constant dialogue and ideas-sharing always lead to uncovering previously unthinkable knowledge. This fortunate sequence of events could not have been possible without the precious support of many remarkable individuals whom I wish to acknowledge next.

I am deeply grateful to my supervisors and co-authors, Maria Montero and Martin Sefton for their thoughtful, valuable guidance and scholarly input. Our countless long meetings have been instrumental not only for the completion of this thesis but also for my personal development as a researcher. I have been constantly impressed by their responsiveness and support in all my endeavors.

The atmosphere within CeDEx[1] has been nothing but extraordinary. I want to thank its members for taking the time to actively engage with my work and provide numerous constructive comments and suggestions. Throughout my PhD experience, I have come to understand the uniqueness of this research group and I am happy to have been part of it. My sincere thanks also go to the administrative staff of CeDEx and of the School of Economics who have always provided me with their excellent assistance.

During the PhD years I have gained many new irreplaceable friendships. My colleagues and friends have been there to make the joyful moments memorable and the tougher times tolerable. I want to thank them for enlarging my cultural and culinary knowledge as well as for all their support throughout.

I also want to thank Theo Offerman who has been an inspirational person in my academic path, providing his valuable support whenever needed. Moreover, I need to praise my family and old friends who, even though from afar, have given me their unconditional support and encouragement.

1 Centre for Decision Research and Experimental Economics

In addition, I gratefully acknowledge the financial support of the Economic and Social Research Council (ESRC), the Network for Integrated Behavioural Science (NIBS) and CeDEx.

Above all, I thank Orestis Kopsacheilis for giving me the energy to approach each challenge with excitement, for opening up new avenues of thought through his inquisitive and insightful character and last but not least, for his tireless research assistance in running the experiments and during our countless engaging discussions.

CONTENTS

1 GENERAL INTRODUCTION 1

2 COMMUNICATION WITH PARTIALLY VERIFIABLE INFORMATION: AN EXPERIMENT 6
- 2.1 Introduction 6
- 2.2 Related literature 8
- 2.3 Partially verifiable information games 9
- 2.4 Experimental design and hypotheses 13
- 2.5 Results 15
  - 2.5.1 Main findings 15
  - 2.5.2 Strategy analysis 17
  - 2.5.3 How can the receiver best protect himself? 26
  - 2.5.4 Comparison of payoffs with the GR equilibrium 31
- 2.6 A different parametrization 31
- 2.7 Conclusion 35
- Bibliography 36

3 CHEAP TALK AND EVIDENCE IN THE LAB 38
- 3.1 Introduction 38
- 3.2 Related literature 39
  - 3.2.1 Cheap talk 40
  - 3.2.2 Disclosure 42
- 3.3 Experimental Design 43
  - 3.3.1 The game 43
  - 3.3.2 Treatments and hypotheses 44
  - 3.3.3 Procedures 45
- 3.4 Results 46
  - 3.4.1 Reporting strategy in CM$_-$ treatments 47
  - 3.4.2 Revelation strategy 49
  - 3.4.3 Acceptance strategy 50
  - 3.4.4 Average payoffs 53
  - 3.4.5 Discussion 54
- 3.5 Follow-up treatments 55
  - 3.5.1 Results 56
  - 3.5.2 Discussion 59
- 3.6 Conclusion 60
- Bibliography 61

4 COMMUNICATION, SKEPTICISM AND LYING AVERSION 63

4.1 Introduction 63
4.2 Related literature 65
4.3 Theoretical analysis 67
4.3.1 The game 68
4.3.2 Equilibrium analysis 69
4.3.3 Introducing lying and sucker costs 70
4.3.4 Hypotheses 77
4.4 Experimental design 79
4.5 Results 80
4.5.1 Buyer behavior 81
4.5.2 Seller behavior 83
4.5.3 Average payoffs 84
4.6 Conclusion 87
Bibliography 90

5 GENERAL CONCLUSIONS 94

A APPENDIX FOR CHAPTER 2 98
A.1 Equilibrium examples 98
A.2 Average payoffs per matching group 100
A.3 Receiver's best response given the empirical distribution in Rv, T=9 101
A.4 Instructions 102
A.4.1 Rv, T=11 102
A.4.2 Sr, T=11 104

B APPENDIX FOR CHAPTER 3 106
B.1 Iterated admissibility 106
B.2 Instructions (No_CM_9) 109

C APPENDIX FOR CHAPTER 4 111
C.1 Instructions 111
C.1.1 CM_23 111
C.1.2 No_CM_23 113
C.2 Seller messages 115
C.2.1 CM_13 treatment 115
C.2.2 CM_23 treatment 119

LIST OF FIGURES

Figure 2.1 Acceptance frequencies for a given value of the highest card 16
Figure 2.2 Distribution of realized vs. reported states in Sr and Rv 18
Figure 2.3 Proportion of bad hands reported as good hands 19
Figure 2.4 Frequency of revealing the higher card (Sr) or verifying the higher message (Rv) when a good hand is reported 20
Figure 2.5 Acceptance and observation frequencies conditional on the value of the verified card (reported good hands only) 22
Figure 2.6 Acceptance frequencies when a good hand is reported and no misreport is observed 23
Figure 2.7 Frequency of good hands given that a good hand was claimed and no misreport was observed 24
Figure 2.8 Acceptance frequencies for a given value of the highest card (T=9) 33
Figure 3.1 Message format 46
Figure 3.2 Frequency with which each combination of values was realized (left) vs. reported (right) 47
Figure 3.3 Proportion of bad hands reported as good 48
Figure 3.4 Percentage of times senders reveal higher card 49
Figure 3.5 Acceptance and observed frequencies conditional on the value of the revealed card 51
Figure 3.6 Distribution of errors across the state space (x-axis: values of the orange card; y-axis: values of the blue card) 52
Figure 3.7 Acceptance frequencies: R_choice vs. CM_9 vs. No_CM_9 57
Figure 3.8 Acceptance frequencies: S_choice vs. CM_9 vs. No_CM_9 59
Figure 4.1 Communication interface 80
Figure 4.2 Frequency with which buyers choose the ex ante optimal action 81
Figure 4.3 Frequency with which buyers choose $a = red$ after $m = red$ conditional on the underlying state 85

LIST OF TABLES

Table 2.1 Payoff Matrix (sender's payoff listed first in each cell) 10
Table 2.2 Acceptance frequencies for pooled data (within 6 categories) 17
Table 2.3 Average payoffs in Sr and Rv 17
Table 2.4 Proportion of misreported hands 18
Table 2.5 Frequencies of different types of reports (bad hands with one high card) 20
Table 2.6 Marginal effects from probit regression 25
Table 2.7 Calculating the best response to the empirical distribution for Rv 27
Table 2.8 Calculating the best response to the empirical distribution for Sr 30
Table 2.9 Average payoffs: Rv vs. GR equilibrium predictions 31
Table 2.10 Average payoffs: Sr vs. GR equilibrium predictions 31
Table 2.11 Acceptance frequencies for pooled data (within 6 categories) 33
Table 2.12 Average payoffs in Sr and Rv when T=9 34
Table 3.1 Payoff Matrix (sender's payoff listed first in each cell) 43
Table 3.2 Main treatments 44
Table 3.3 Acceptance frequencies across treatments 53
Table 3.4 Average payoffs: CM_ vs. No_CM 54
Table 3.5 Average payoffs: R_choice, No_CM_9, CM_9 56
Table 3.6 Average payoffs: S_choice, No_CM_9, CM_9 58
Table 4.1 Payoff Matrix (seller's payoff listed first in each cell) 68
Table 4.2 Treatments summary 79
Table 4.3 No. of observed red and blue states across treatments 80
Table 4.4 Distribution of buyer actions conditional on the message received 81
Table 4.6 Message types 83
Table 4.7 Message frequencies 83
Table 4.8 Message distribution across states (cells depict overall counts) 83
Table 4.9 Average payoffs across treatments 85
Table 4.10 Average payoffs across treatments conditional on the underlying state 86
Table 4.11 Average payoffs across treatments 86

Table 4.12 Distribution of errors 87
Table A.1 Sender strategy for receiver's worst equilibrium when T=11 99
Table A.2 Sender strategy for receiver's worst equilibrium when T=9 99
Table A.3 Average payoffs per matching group: T=11 100
Table A.4 Average payoffs per matching group: T=9 100
Table A.5 Calculating the best response to the empirical distribution for Rv, T=9 101

1

GENERAL INTRODUCTION

Communication about private information permeates our economic interactions. Although its volume might seem overwhelming in today's Information Age, "knowledge"-sharing is in fact a paramount efficiency of competitive markets as Hayek (1945) argued in his seminal essay. Of course, not all information sharing is actionable or has economic implications. Job interviews, financial investments, voting decisions - these are all examples of situations in which one party is interested in sharing all or some of their private information to influence the actions of another party. Such private information sharing is strategic - the overarching theme of this thesis.

There are several ways in which private information can be shared strategically. For example, Spence (1973) showed that a costly action such as acquiring formal education can be used to signal the individual's high ability. Another possibility is to simply engage in unverifiable and non-costly (cheap) talk or in non-costly but verifiable communication. This thesis focuses on the strategic transmission of information through non-costly actions.

Research on this subject goes back at least to 1980 when Green and Stokey (1980) introduced a theoretical model for the analysis of unverifiable messages about private information. Soon after, Crawford and Sobel (1982) discuss a similar model which has now become the foundation of the "cheap talk" research. The typical framework in which this has been analyzed is that of a sender-receiver game. The findings of this research suggest that the privately informed senders would prefer to act dishonestly as long as it is in their interest to do so.

Notwithstanding, when communication is verifiable, the seminal paper of Milgrom and Roberts (1986) proved that in a setting where senders can only tell the truth or hide it without being dishonest, there is never an incentive for senders to hide information. An under-explored area of research is that in which senders and receivers interact in environments with characteristics from both verifiable and unverifiable non-costly communication settings. The present thesis delves into this topic in the next three chapters which are briefly summarized below. Although these can be read independently, they are encompassed by common themes and research methodologies.

In chapter 2 we look at a setting where both forms of information - verifiable and unverifiable - are simultaneously present. Specifically, we are interested

in the effect of verification control when information is partially-verifiable. In researching this topic we make use of a laboratory experiment based on a persuasion game introduced by Glazer and Rubinstein (2004). The game involves a sender, perfectly informed about a bi-dimensional state, who is trying to convince the uninformed receiver to take an action in the sender's interest. In doing so, the sender sends a message (possibly untruthful) to the receiver about the values of the two state-dimensions; however, only one dimension is verifiable. We compare strategic behavior in the case where the receiver decides which dimension to verify with that where the sender has verification control. Theoretically, the two environments are equivalent as the receiver's best equilibrium leads to the same maximum payoff (Glazer and Rubinstein, 2004, 2006). However, the equilibrium strategies when the sender is in control are very compelling: the sender reveals the dimension with the highest value, and the receiver accepts if the value observed is sufficiently high.

Our experimental data shows that when the sender has verification control he reveals the dimension with the highest value as predicted. The receiver's payoff is though reduced due to noise in his acceptance strategy. Both players' behavior is noisy when the receiver has verification control, but not all deviations from the theoretical prediction hurt the receiver. As a result, average payoffs are not significantly different between treatments. A robust finding of this chapter is that receivers are more skeptical about the observed evidence when senders have verification control compared to the case where receivers do so. Skepticism in this context means that the receiver behaves as if distrusting that the state is the one in which his optimal action is also the sender's preferred one. Furthermore, we investigate under which circumstances the receiver can best protect himself from deceptive behavior. We show that the receiver's best reply to sender's observed behavior in our experiment is close to the optimal strategy identified by Glazer and Rubinstein.

In chapter 3, we focus on the setting in which the sender has verification control. In this investigation we ask how does cheap talk influence receivers' perception of the partially-informative evidence revealed by the sender. Could the noise in receiver's behavior observed in the previous chapter be due to non-verifiable communication? We address these questions by using a similar methodology as in chapter 2. Specifically, we compare behavior in the setting where the sender has verification control with a similar one where only non-costly, verifiable communication is possible as the option of cheap talk is removed.

Our results show that cheap talk messages prime more mistakes in the appraisal of the evidence provided. In particular, these messages make receivers

more skeptical about favorable[1] evidence and less so about unfavorable information, leading to a negative impact on receivers' payoff. We find this surprising as receivers can always ignore cheap talk.

To eliminate possible design aspects that might lead to receivers feeling compelled to account for messages in their decisions, we run two follow-up treatments. These treatments are designed to clarify to receivers that messages need not be accounted for in their decision making. In the first treatment receivers can actively ignore the message by choosing not to have it displayed on their screens. In the second treatment senders can choose not to send the message. In these follow-up treatments we implicitly vary the amount of control receivers have over the unverifiable communication. We find that the negative effect of cheap talk is mitigated by increasing receiver's control over the communication space. This chapter sheds light on how individuals differentiate between more and less relevant types of information, highlighting that communication can make the involved agents worse off.

Skepticism stemming from communication has been previously documented by Ert et al. (2014). It is though unclear from their study what potential mechanism could give rise to this effect. Building on this, the fourth chapter explores whether the skepticism effect of cheap talk is due to receivers' aversion to being deceived. We refer to this as "sucker aversion". Using a simple theoretical model we show that a dislike for taking an action favorable for the sender when the sender is being deceptive can lead to skepticism and further on, to negative payoff consequences. On the other hand, senders' lying aversion (Abeler et al., 2016) can counter this effect and lead to material gains. We identify settings in which one of the two effects dominates. Notably, when players' interests are more likely to be aligned, the "sucker aversion" effect takes over. Conversely, when the interests are less likely aligned, the lying aversion effect dominates. We test these theoretical predictions using a laboratory experiment.

As opposed to the previous chapters, we abstract from settings where verifiable communication is available. This is so that we can isolate the effect of communication more clearly. We focus on a one-shot seller-buyer environment with asymmetric information about the state of the world. In this game, the seller has perfect information about the true state but the buyer does not. There are two possible states, blue and red, and the realized state is randomly determined by drawing a ball from an urn containing three balls of red and blue colors. Which of the two colors is more frequent depends on the treatment. The seller would like the buyer to always buy the red product. However, the buyer prefers to buy the red product only if a red ball was drawn, and

1 The favorable-unfavorable dichotomy stems from the posterior probability that the state is such that the receiver would like to take the sender's preferred action (favorable) or not (unfavorable).

the blue product if a blue ball was drawn. We vary two dimensions: whether the seller is allowed to send a message to the buyer and whether the optimal strategy for the buyer is to buy the blue or the red product.

We find that when interests are more likely aligned (the urn contains two red balls and a blue one), communication increases skepticism but not significantly so. Consequently, there is no significant negative effect on players' payoffs which is evidence against the sucker aversion hypothesis. In contrast, when interests are more likely opposed (the urn contains one red ball and two blue ones), communication reduces skepticism resulting in a (weakly) significantly positive effect on the seller's payoff without affecting the buyer. This is though not due to the hypothesized lying aversion effect; rather, the explanation is more likely related to buyers' concern for efficiency.

BIBLIOGRAPHY

Abeler, J., Nosenzo, D., and Raymond, C. (2016). Preferences for truth-telling. *CESifo Working Paper Series No. 6087*.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Ert, E., Creary, S., and Bazerman, M. (2014). Cynicism in negotiation: When communication increases buyers' skepticism. *Judgment and Decision Making*, 9:191–198.

Glazer, J. and Rubinstein, A. (2004). On optimal rules of persuasion. *Econometrica*, 72(6):1715–1736.

Glazer, J. and Rubinstein, A. (2006). A study in the pragmatics of persuasion. *Theoretical Economics*, 1:395–410.

Green, J. R. and Stokey, N. L. (1980). *A two-person game of information transmission*. Harvard Institute of Economic Research.

Hayek, F. (1945). The use of knowledge in society. *American Economic Review*, 35:519–30.

Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821.

Spence, M. A. (1973). Job market signaling. *Quarterly Journal of Economics*, 87:355–374.

2

COMMUNICATION WITH PARTIALLY VERIFIABLE INFORMATION: AN EXPERIMENT

*

2.1 INTRODUCTION

We report experiments based on persuasion games in which an informed agent (sender) wishes to persuade an uninformed agent (receiver) to take a certain action, and we ask how the receiver can best protect himself against false claims by the sender. Such situations are commonplace. For example, in a job interview an applicant wishes to persuade an employer to hire them; in a car showroom a salesman wishes to persuade a customer to purchase a car. Also, in such situations the uninformed party is often able to gather some hard evidence to assist making a good decision, e.g. by testing the applicant or test-driving the car, but typically it is infeasible or prohibitively costly to eliminate all asymmetric information. Thus, we focus on games in which the sender's information is only partially verifiable. These games are based on theoretical models of Glazer and Rubinstein (2004, 2006) and we bring experimental evidence to bear on them.

In the games we study the players' interests are aligned in some states of the world but opposed in others, where the state of the world is based on the values of two aspects. These values are known to the sender, but not to the receiver, and the sender sends a message about the two aspects. In our first game the receiver can then choose to observe one of the aspects before taking an action (we refer to this as the "Receiver Verifies" game). This game is based on the model introduced in Glazer and Rubinstein (2004), who identify a receiver's optimal strategy assuming that the receiver is able to commit to a strategy. The game we implement is an extensive-form game in which the receiver cannot commit. Nature informs the sender of the value of the two aspects, the sender then announces (possibly untruthfully) the values to the receiver, the receiver then chooses an aspect to observe and then takes an action. This game has a sequential equilibrium that gives the same payoff to the receiver as the optimal mechanism (Glazer and Rubinstein, 2004, section 7), and we use this as a prediction (the "GR" prediction) for how the game will be

* This chapter is based on joint work with Maria Montero and Martin Sefton.

played. Essentially, the receiver verifies the aspect with the higher announced value and accepts if the observed value is sufficiently high. However, we note that there are many other sequential equilibria giving a lower payoff to the receiver, and no obvious reason why the receiver's preferred equilibrium would be played.

The second game we study differs from the first in that the sender decides which aspect will be observed by the receiver. This "Sender Reveals" game is based on the model introduced in Glazer and Rubinstein (2006). Using a mechanism design approach, they show that the receiver's optimal strategy (with commitment) yields the same payoff as in the model where the receiver verifies. Also in this case, our laboratory extensive form game has a sequential equilibrium that implements the receiver's optimal payoff. This time, the equilibrium that is optimal for the receiver is particularly simple and intuitive: the sender reveals the aspect with the highest value and, realizing this, the receiver decides accordingly.

In summary, each game has multiple equilibria, and the maximum equilibrium payoff for the receiver is the same in both games. Since there are multiple equilibria in both games, equilibrium theory (absent any refinement criteria) does not offer a sharp prediction about which game is better for the receiver. However, because the equilibrium supporting the receivers' optimal payoff is particularly simple and intuitive in Sender Reveals, we expect behavior to be close to this equilibrium in Sender Reveals. For Receiver Verifies, we find the corresponding equilibrium less compelling, and hence we hypothesize that observed behavior will be further from the GR prediction.

The main contribution of this chapter is the empirical study of behavior in settings of strategic information transmission in which both cheap and hard information are present. Moreover, we compare settings in which either the receiver or the sender have control over the hard information set. As expected, we find that outcomes are closer to the GR prediction in the Sender Reveals game. Nevertheless, there are deviations in both games. While there are fewer deviations in the Sender Reveals game, these systematically reduce receivers' payoffs. While sender behavior is very close to that predicted, receivers' acceptance strategy is noisy, resulting in payoffs for the receiver significantly below the GR prediction. In the Receiver Verifies game there are more deviations from the GR prediction, but not all deviations are detrimental to the receiver. The net effect is that receiver's payoffs are similar in the two games.

Our experiment also allows us to address the question of how the receiver can best protect himself against false claims, given the actual behavior of senders. As we anticipated, the receiver should play according to the GR prediction in Sender Reveals. Since sender's behavior is very close to the theoretical prediction, the receiver's best response would be to follow the theoretical

prediction as well. More surprisingly, despite the noisy behavior of the sender in Receiver Verifies, the GR strategy is close to being optimal for the receiver also in this setting.

The chapter is organized as follows. In the next section we discuss the related literature. In Section 3 we introduce the games and the theoretical predictions. Section 4 presents the experimental design and hypotheses. Section 5 describes the results. Section 6 analyses an alternative parametrization of the game and Section 7 concludes with a discussion.

2.2 RELATED LITERATURE

Our experiment studies strategic communication in a setting where an informed sender can send messages about private information and these messages are partially verifiable. Specifically, we base our set-up on models introduced in Glazer and Rubinstein (2004, 2006).

Glazer and Rubinstein (2004) analyze a situation where a sender is privately informed about a multi-dimensional state of the world and can send a message about this to a receiver. The receiver can then choose a single dimension of the state to observe, and can thus verify part of the message, before taking one of two actions, Accept or Reject. The sender prefers the receiver to accept independent of the state, whereas the receiver's optimal action depends on the state. The authors focus on finding optimal mechanisms from the receiver's point of view, i.e. mechanisms that maximize the receiver's expected payoff[1].

Glazer and Rubinstein (2006) modify this model by removing the receiver's option to verify and instead allowing the sender to reveal truthfully one dimension of the state. They show that the optimal mechanism (from the receiver's point of view) in this case yields the same expected payoff to the receiver. In this sense the receiver does not suffer by losing verification control.

These settings combine elements of cheap talk games and verifiable message games. Few studies can be included in this category: Lipman and Seppi (1995) examine the role of competition among senders in a model where information is partially verifiable while Forges and Koessler (2005) characterize the equilibrium set of such games when a communication mediator is present.

1 In the optimal mechanism the receiver sometimes makes errors in the sense that there are states where they accept but would have been better off rejecting, and vice versa. See Carroll and Egorov (2017) for a related setting where the receiver is able to eliminate errors altogether.

The literature on Bayesian persuasion games is related to the models we study. These are sender-receiver games where the sender wants to persuade the receiver without the use of cheap talk by partially or fully certifying his type. The focus of this body of work is on the analysis of information certification strategies that maximize sender's persuasiveness. The canonical model has become that of Kamenica and Gentzkow (2011) but the setting has been analyzed also in conditions with multiple receivers (Alonso and Câmara, 2016) or repeated interaction (Forges and Koessler, 2008). We differ from these settings by allowing cheap talk and focusing on receiver's interests.

A substantial experimental literature has examined cheap talk games based on the model of Crawford and Sobel (1982) (see Blume et al., 2017 for a review), and similarly previous experiments have studied verifiable message games based on Milgrom and Roberts (1986) model (see e.g. Forsythe et al., 1989; Hagenbach and Perez-Richet, 2018; Jin et al., 2015; King and Wallin, 1991; Li and Schipper, 2018; Penczynski and Zhang, 2016). The game we use differs from a cheap talk game in that messages can be (partially) backed by evidence, and differs from a verifiable message game in that we force partial disclosure (whereas in verifiable message games senders may remain silent or fully reveal their type). This combination has important theoretical implications. For example, the typical theoretical result in verifiable message games is that senders always reveal their type (the unravelling principle). This, by design, does not apply to games with partial disclosure such as ours. As far as we are aware, ours is the first experimental study of a setting that combines cheap talk and partially verifiable message games.

2.3 PARTIALLY VERIFIABLE INFORMATION GAMES

The Receiver Verifies game (henceforth Rv) is based on Glazer and Rubinstein (2004)[2]. The sender is dealt 2 cards, one orange and one blue, each taking a value between 1 and 9. Each combination of cards represents a state of the world and is equally likely. The sum of the two cards determines whether the state is good (the sum is at least T) or bad (the sum is lower than T). We focus on the cases where this threshold is equal to 11. After observing the two cards, the sender sends a (possibly untruthful) message reporting the values of the two cards[3]. After observing the message, the receiver chooses a card to observe. Next, upon observing the true value of the selected card, the receiver decides whether to accept or reject. The sender prefers the receiver

2 Our exact parametrization is taken from the online experiment of Glazer and Rubinstein at http://gametheory.tau.ac.il/exp5/. In this experiment, the receiver is a computerized player playing an undisclosed strategy.

3 Hence, communication must be precise; the sender cannot leave the message blank or send a vague message (cf. Serra-Garcia et al., 2011).

to accept, irrespective of the state of the world, whereas the receiver would like to accept only when the state is good and reject otherwise. The interests of the two players are aligned if the hand is good and opposed if the hand is bad. More specifically, payoffs are as given in table 2.1.

Table 2.1: Payoff Matrix (sender's payoff listed first in each cell)

	Receiver accepts	Receiver rejects
Sender has a good hand	(1, 1)	(0, 0)
Sender has a bad hand	(1, 0)	(0, 1)

The Sender Reveals game (henceforth Sr) differs from the Rv game only in that, after the message is sent, it is the sender who chooses which card is observed by the receiver. For the equilibrium analysis it is convenient to begin with the Sr game.

Sequential equilibria of Sr

The Sr game has many sequential equilibria. Among these, one of them stands out as particularly compelling because the sender's strategy is simple and intuitive: always display the card with the highest value. The receiver then accepts if, taking into account that the value observed is the highest of the two, the hand is more likely to be good than bad. This results in the receiver accepting if the observed value is at least 7. Note that this equilibrium makes deterministic predictions about which hands are accepted: senders who have a high card (7+) are accepted for sure, while other senders are rejected. This results in some bad hands (such as (7,1)) being accepted and some good hands (such as (6,5)) being rejected.

Glazer and Rubinstein (2006) show that in an Sr type of game without messages the same outcomes are obtained in a sequential equilibrium which yields the highest expected payoff to the receiver. There, the sender simply reveals his best card and the receiver accepts if this is 7+[4]. We take this equilibrium as the basis for our hypotheses[5].

GR prediction (Sr). *The sender always reveals the higher of the two cards. The receiver accepts if the observed value is at least 7.*

Messages play no role in equilibrium in the sense that, holding the value and color of the observed card constant, if the receiver behaves differently for messages m and m' then the less favorable message will not be observed on

4 Their result is even stronger: the receiver would not be able to do better even if he were able to commit to any strategy.

5 This equilibrium also survives iterative elimination of weakly dominated strategies in a game where the sender is restricted to simply revealing a card (without an accompanying cheap talk message). See Appendix B.1 for details.

the equilibrium path. Hence, in equilibrium the receiver behaves as if ignoring the message[6]. Because of this, we have not included any statement about messages in the theoretical prediction. Alternatively, the theoretical prediction could involve a requirement that the sender reports a good hand (either always or when he has one high card) and the receiver rejects if the sender reports a bad hand[7]. There are also equilibria which give rise to different outcomes and a lower payoff for the receiver[8].

Sequential equilibria of Rv

For the Rv game there is a family of equilibria that are equivalent to the one we have selected for Sr. They are equivalent in the sense that all senders who have a 7+ card are accepted for sure, and all senders who do not are rejected for sure. However, multiple equilibrium strategies can support this outcome. Perhaps the most natural is that in which the sender truthfully reports good hands that include a 7+ card. If the hand is bad but the highest card is 7+, the sender inflates the lower card in such a way that the reported values for the two cards add up to at least 11 while keeping the highest report truthful (for example, (7,1) may report (7,5))[9]. The receiver checks the highest message and accepts if it is true (and the value is 7+)[10].

We take the following as the theoretical prediction.

GR prediction (Rv). *The sender reports good hands truthfully, and inflates bad hands if one of the cards is 7+. If the sender reports a bad hand, the receiver rejects. If the sender reports a good hand, the receiver checks the higher of the two messages, and accepts if the observed value coincides with the message and is at least 7.*

Because there are many strategies compatible with the GR outcome, we do not specify complete strategies in the equilibrium prediction. Instead, we

6 In other words, communication may be informative in equilibrium (i.e., the receiver's beliefs may be affected by the message) but cannot be persuasive (the receiver's optimal action cannot change as a result of the message); see Lipnowski and Ravid (2017). We return to this point in our analysis of messages.

7 We could also add a requirement that the message does not contradict the displayed card in order for the receiver to accept. Glazer and Rubinstein (2004) impose this requirement in their mechanism.

8 See Appendix A.1 for an equilibrium that gives the receiver a lower expected payoff.

9 Note that we have not described a complete strategy for the sender; we return to this below.

10 There are many other equilibrium strategies that lead to the same payoffs. For example, it is not essential that the receiver always checks the highest card. The receiver may check a card at random if both reports are 7+; the sender then takes care not to inflate the lower card above these values if the hand is bad. There are also equilibria where messages are not interpreted literally; for example, there are equilibria where only two (arbitrary) messages are used, and the receiver checks the blue card after one message and the orange card after the other. For our analysis of the experiment we will assume that subjects intend the messages to have a literal meaning.

focus on features of the strategies that are common to all or most of these equilibria[11].

There are many other sequential equilibria, including cases where messages are not informative (babbling) and equilibria that are better for the sender. Glazer and Rubinstein (2004) show that no equilibrium gives the receiver an outcome better than the GR prediction. Hence, being able to choose a dimension to verify does not confer any advantage to the receiver according to theory.

Furthermore, while displaying the highest card in Sr is very natural, it is not obvious that the receiver should always check the highest message, or should accept whenever the observed value is 7+. Suppose the receiver checks the highest message and accepts if the message is true and the value is 7+. Glazer and Rubinstein (2004) show that there exists a best response for the sender which leads to a sequential equilibrium. However, there are other best responses for the sender that do not lead to a sequential equilibrium. That is, the GR family of equilibria require that the sender breaks ties in a specific way.

As an illustration, suppose the message is (7,4). The GR equilibrium requires the receiver to check the first card and accept if it is 7. However, if the sender only inflates bad hands up to the bare minimum needed to make up a good hand, after receiving message (7,4) and verifying that the 7 is correct, the hand is equally likely to be (7,4), (7,3), (7,2) and (7,1). Only one out of four is a good hand, hence the receiver should reject. In order to have an equilibrium, the sender cannot concentrate the lies on messages that add up exactly to the threshold; one possibility is that (7,3) reports (7,6), (7,2) reports (7,5) and (7,1) reports (7,4).

Another issue is whether the receiver should always verify the card with the highest reported value (or even a card with a 7+ reported value). If the sender is more likely to be lying about the lowest reported value, it would be optimal for the receiver to check the lowest report instead.

The GR equilibrium also requires the receiver to reject some good hands, such as (6,6). If (6,6) is the only hand that sends message (6,6), then the receiver would know for sure that this is a good hand and should accept instead. In order to obtain the GR equilibrium, one needs to assume that either (6,6) reports one of the cards as 7+ (even though this will be discovered for sure given the strategy of the receiver) or there are bad hands that also report (6,6), so that (6,6) is not unambiguously a good hand. In both cases, there are sender types that are lying even though they have nothing to gain from doing so. A

11 The coincidence of the observed value with the message is not essential to the equilibrium, but it is explicitly required in the optimal mechanism of Glazer and Rubinstein (2004).

lexicographic preference for telling the truth (i.e., players only lie if they have something to gain) would eliminate this equilibrium[12].

In contrast with Sr, messages play an important role in Rv: they guide the receiver in his choice of which card to check[13,14].

GR prediction (focusing on outcomes, both treatments). *Hands with one high card (7+) are always accepted, while hands without a high card are always rejected.*

Since the GR prediction is more compelling in Sr as argued above, we hypothesize behavior will be closer to predictions in Sr. We also expect the receiver to achieve a higher expected payoff in Sr, since the GR prediction gives the receiver his maximum possible equilibrium payoff. The next section describes our experiment and formally states the hypotheses it is designed to test.

2.4 EXPERIMENTAL DESIGN AND HYPOTHESES

The experiment was conducted in the CeDEx laboratory at the University of Nottingham, UK. There were 192 subjects, recruited from a university-wide pool of undergraduate and graduate students using ORSEE (Greiner, 2015). The experiment was programmed in z-Tree (Fischbacher, 2007).

Our experiment varies the game (Sr and Rv) across sessions. 4 sessions of each treatment were conducted with 24 subjects per session, and each session was divided into two matching groups. This gives us 8 independent observations per treatment. 65% of participants in the Sr treatment and 64% in the Rv one were female.

Upon arrival at a session, subjects were randomly allocated a seat number and given a set of instructions, which were then read out loud by the experimenter[15]. The decision-making part of the experiment consisted of 30 rounds,

12 A similar refinement has recently been used by Hart et al. (2017).

13 Glazer and Rubinstein (2006) did not include cheap-talk communication in their Sr game. We have done so in our experiment for the sake of comparability between treatments.

14 There are also babbling equilibria where messages are not informative (for example, all sender types send each of the possible messages with equal probability, or all sender types send the same message). The receiver does not condition the checking decision on the message (for example, the receiver always checks the same card, or checks each of the two cards with equal probability). Upon observing the true value of one of the cards, the receiver takes the optimal acceptance decision given that the other card is equally likely to be any value between 1 and 9. The receiver then accepts if and only if the observed card is at least 6 (since in 5/9 of cases the sum is at least 11).

15 See Appendix A.4 for a copy of the instructions.

where in each round, subjects were randomly matched to play the relevant game. Subjects were re-matched within their matching group at the beginning of each round, but retained the same role (sender or receiver) during the entire session.

At the beginning of each round, each sender is dealt two cards (blue and orange). Each of the cards is equally likely to be any integer value between 1 and 9, and all draws are independent across colors and senders[16]. A hand is defined as "good" if the sum of the values of the two cards is at least 11. Having observed the two cards, the sender sends a message to the receiver of the form "The value of the orange card is __; The value of the blue card is __". Next, the receiver chooses one of the cards to observe (Rv treatment) or the sender chooses one of the cards for the receiver to observe (Sr treatment). The receiver then accepts or rejects. The sender earns 1 point if the receiver accepts and 0 if the receiver rejects; the receiver earns 1 point if he accepts a good hand or rejects a bad hand, and 0 otherwise.

At the end of each round, a summary screen displayed the true values of the two cards, the message sent by the sender, the card chosen to be observed (by the receiver or by the sender, depending on treatment), the receiver's decision, and the point-earnings of the two subjects.

At the end of the experiment subjects received their accumulated earnings for the 30 rounds (1 point = £0.50) plus a £3 participation fee. There was a total of 1440 independent draws for each treatment out of which 631 were good hands. Each session lasted around 90 minutes, and average earnings were £13.

With this design, we are able to test whether behavior conforms more closely to GR predictions in Sr or Rv, and whether the receiver's payoff is higher in Sr or Rv. GR predictions are deterministic and identical in both treatments: any given hand will be either accepted or rejected. Hence, we use a hit rate metric to test conformance of behavior with predictions.

Specifically, we count how often the receiver's actual decision matches the predicted decision[17]. Our first Hypothesis is:

16 For the Rv treatment the card draws were randomized using the random number generator during the session. To enhance comparability across games, we then used these realizations in the corresponding sessions of the Sr treatment. This allows us to perform the statistical comparisons on paired observations. Unless otherwise specified, all statistical tests are two-tailed signed-rank tests at the 5% significance level and taking the matching group as the unit of observation.

17 This metric has the advantage of summarizing the closeness to the prediction by a single number. A disadvantage is that it potentially overestimates the compliance with equilibrium predictions, since an equilibrium outcome may arise from non-equilibrium strategies. We discuss the closeness of the strategies to equilibrium predictions as well.

Hypothesis 1. *(Behavior closer to GR prediction in Sr)* The proportion of GR outcomes is higher in Sr than in Rv.

Since the GR prediction yields the maximum equilibrium payoff for the receiver, we also hypothesize:

Hypothesis 2. The receiver achieves a higher expected payoff in Sr than in Rv.

2.5 RESULTS

2.5.1 *Main findings*

The theoretical prediction is sharp: hands that contain a sufficiently high card are always accepted, while other hands are rejected. Below we report the observed acceptance frequencies conditional on the value of the higher of the two cards, for good hands and bad hands separately. The dashed lines in figure 2.1 represent the theoretical prediction. Note that, for all hands, the acceptance frequency is also the sender's average payoff; the receiver's average payoff is the same as the sender's payoff for good hands, and one minus the sender's payoff for bad hands. Table 2.2 is the numerical companion of figure 2.1, making it easier to translate the effect of the acceptance strategies across games to average payoffs. While figure 2.1 presents the data in a disaggregate manner, in table 2.2 we pool the hands for the relevant categories.

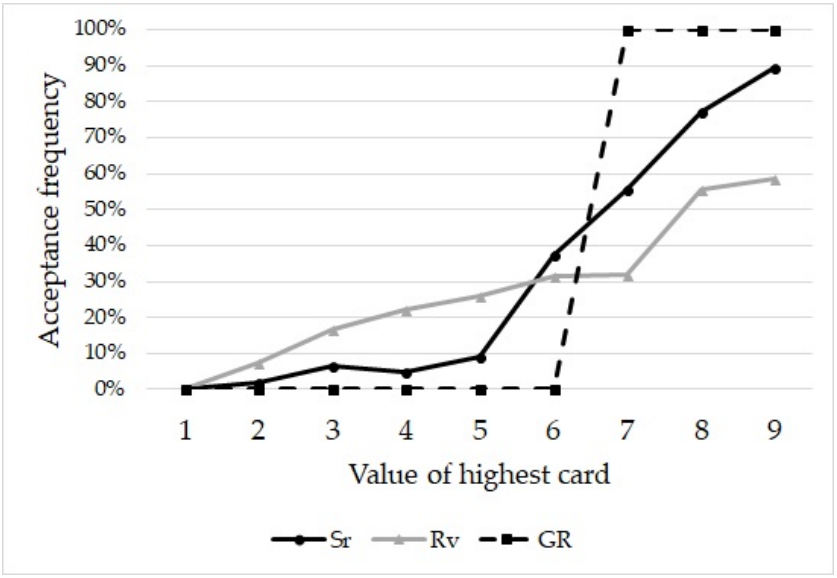
We observe a difference in behavior, but not in average payoffs. Looking at figure 2.1, it is clear that although there are deviations in both treatments, outcomes are closer to the equilibrium prediction in Sr for bad hands, while there is very little difference for good hands. This is confirmed by a comparison of hit rates: overall, the GR equilibrium predicts the outcome of 82.64% of interactions in Sr and 75.07% of interactions in Rv ($p - value = 0.008$). Our results thus confirm Hypothesis 1.

Result 1. *The proportion of GR outcomes is higher in Sr than in Rv.*

This finding is driven by a difference in the hit rates for bad hands (81.21% of interactions in Sr and 68.23% of interactions in Rv have the predicted outcome). The information in panel (a) in table 2.2 supports this: the acceptance frequencies in Sr for the two types of bad hands are closer to the GR prediction, and significantly different than the frequencies in Rv. For good hands, there is no significant difference (84.47% in Sr versus 83.84% in Rv have the predicted outcome). For this case, panel (b) in table 2.2 shows that the acceptance frequencies are different only when the hand contains a 6, though even

then, Sr behavior is not as close to equilibrium predictions as in the case of a bad hand.

(a) Bad hands



(b) Good hands

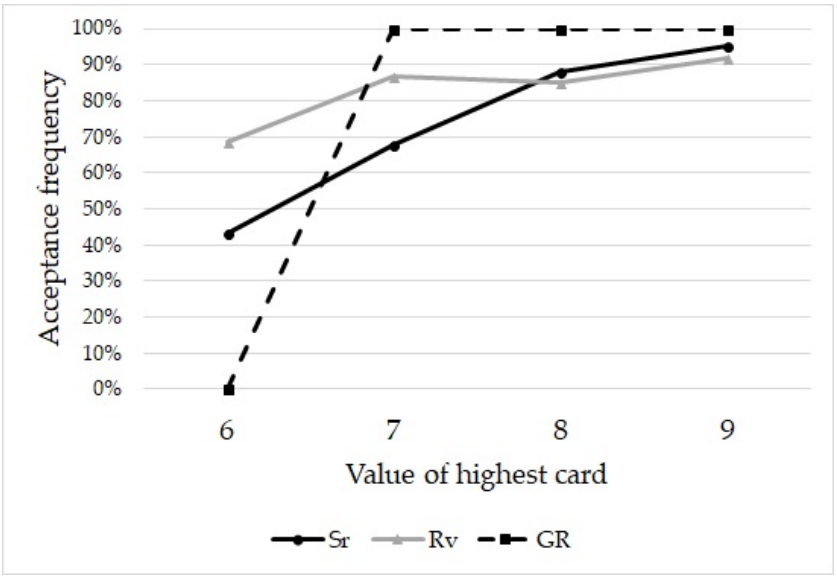


Figure 2.1: Acceptance frequencies for a given value of the highest card

Table 2.2: Acceptance frequencies for pooled data (within 6 categories)

Type of hand		Freq. GR	Freq. Sr	Freq. Rv	count	p-value
(a) Bad hands	No high card (1-6)	0.000	0.139	0.228	591	0.039
	One high card (7-9)	1.000	0.679	0.440	218	0.039
	All bad hands	0.270	0.284	0.286	809	1.000
(b) Good hands	No high card (1-6)	0.000	0.431	0.686	51	0.016
	One high card (7-9)	1.000	0.869	0.885	580	0.468
	All good hands	0.919	0.834	0.869	631	0.250

Note: Counts are for each treatment.

As for payoffs, figure 2.1 and table 2.2 suggest no significant difference, since the overall acceptance frequencies are very similar for the two treatments. Formal testing confirms this (see table 2.3).

Table 2.3: Average payoffs in Sr and Rv

	Sr	Rv	p-value
Sender	0.525	0.541	0.468
Receiver	0.767	0.782	0.641

Result 2. *Average payoffs do not differ significantly between treatments.*

In the following section, we analyze in more detail the differences in behavior between treatments.

2.5.2 *Strategy analysis*

2.5.2.1 *Senders' reporting strategy*

Figure 2.2 presents the distribution of the realized and the reported hands. The middle panel (b) contains the distribution of the values observed by senders. Recall that the random draws are identical across treatments. There are 1,440 random hands in each of the treatments. As each combination of values was equally likely, the distribution of the realized draws is approximately uniform. The left side plot (panel (a)) presents the distribution of the values reported by senders in the Sr treatment; the right side plot (panel (c)) represents the distribution of reports in the Rv treatment.

Result 3. *Senders inflate bad hands in their reports in both treatments.*

If senders were always truthful, there would be no difference between the realized and the reported distributions. Instead, looking at the (a) and (c) panels we observe a shift in the distribution of reported values towards the area

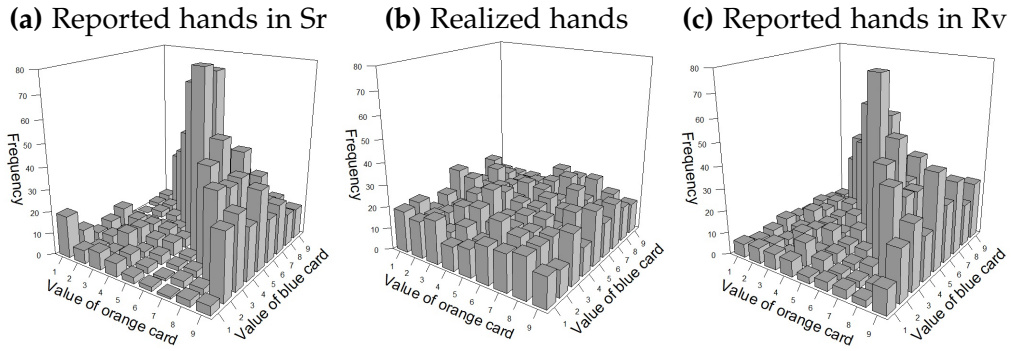


Figure 2.2: Distribution of realized vs. reported states in Sr and Rv

where these add up to at least 11. From table 2.4 we note that the overall reporting patterns are similar in both treatments: around 80% of bad hands are misreported (nearly always as good hands), and a small proportion of good hands are misreported as well (nearly always as good hands, and usually inflated).

Table 2.4: Proportion of misreported hands

	Sr	Rv	p-value
Overall	0.517	0.490	0.461
Good Hands	0.152	0.082	0.016
Bad Hands	0.801	0.807	0.813

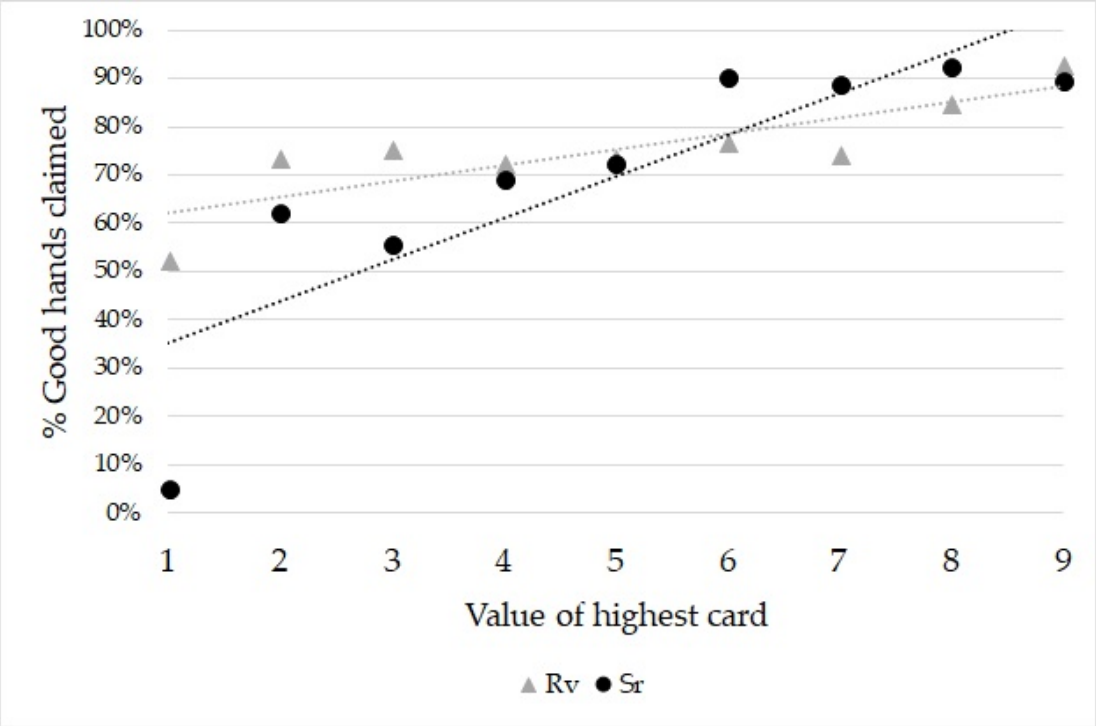
We also note that very few subjects can be characterized as truth-tellers (never misreporting). Out of 48 senders per treatment, we observed only 2 truth-tellers in the Rv treatment and 1 in the Sr treatment[18].

Even though the overall frequency of misreports is similar in Sr and Rv, there are some differences in how these reports are distributed and in the type of misreports used across treatments. Focusing on bad hands, figure 2.3 presents the frequency with which the sender reports a good hand for each possible value of the highest of the two cards. This proportion is relatively flat in Rv, where either of the two messages can be checked, but tends to increase with the value of the highest card in Sr[19]. The proportion of senders who own

18 Thus, our setting does not induce a very strong norm of honesty (cf. Abeler et al. (2016)). We conjecture that this is due to the conflict of interests which can give rise to envious preferences shown to crowd out lying aversion (Cabrales et al., 2016; Minozzi and Woon, 2013)

19 The plausibility of a lie in Sr depends heavily on the value of the highest of the two cards. For example, a sender with cards (4,2) would have to display the 4 and report the non-displayed card to be at least 7 (or, worse, display the 2 and report the other card to be 9), while a sender with cards (1,7) can display the 7 and report the other card to be at least 4. Intuitively, receivers should be skeptical in the first case (since the sender claims to be displaying the lower of the two cards) but not necessarily in the second. In contrast, in Rv, both (4,2) and (1,7) could send message (4,7), and the lie may potentially be discovered or go unnoticed

up to having a bad hand is similar in both treatments, but such reports tend to be more useful for the receiver in Rv (in Sr they are more concentrated on hands where both cards are very low and the receiver would be very likely to guess a bad hand from the observed card alone).



Note: 809 counts for each treatment

Figure 2.3: Proportion of bad hands reported as good hands

Recall that, from the point of view of the GR theoretical prediction, there are two types of bad hands: hands that contain one high card (7+) and hands that do not. The first ones are predicted to always be accepted (after the receiver observes the highest of the two cards), while the latter ones are predicted to always be rejected (irrespective of which card the receiver observes). The theoretical prediction is more precise for bad hands that involve one high card: it is expected that the sender reports a good hand by inflating the value of the lower card only, while keeping it below the value of the higher card. For example, a sender with cards (7,2) may report (7,5) but not (9,2) or (7,8) in a GR equilibrium. The latter reports would result in the hand being rejected if the receiver checks the higher message as predicted.

From table 2.5 we note that overall, for bad hands that would be accepted in equilibrium, the proportion of senders that comply with the theoretical prediction is low in Rv: 61% of senders report a good hand while keeping the highest of the two reports truthful, compared with 79% in Sr. This difference is significant ($p - value = 0.031$). The counterpart result is that senders in Rv

in each case depending on the card checked by the receiver (under the GR prediction the receiver would always check the 7, but not all receivers behave this way).

Table 2.5: Frequencies of different types of reports (bad hands with one high card)

	Sr	Rv
Truthful reporting	7%	17%
Misreport in lower message only	79%	61%
Other misreports	14%	22%

are more truthful than in Sr or misreport both values more often, actions that unequivocally help the receiver.

Result 4. *Senders' reporting strategy is closer to the theoretical prediction in Sr. In Rv, the sender deviates in ways that may help the receiver to reject bad hands.*

2.5.2.2 *Verification/revelation strategy*

When the sender reports a bad hand, the verification strategy is immaterial since presumably the sender is telling the truth (indeed, more than 98% of all reported bad hands are actually bad hands in both treatments). Focusing on cases where the sender reports a good hand, figure 2.4 presents the frequency with which the sender reveals the higher card (Sr) or the receiver verifies the higher message (Rv); cases where both cards/messages are equal are excluded.

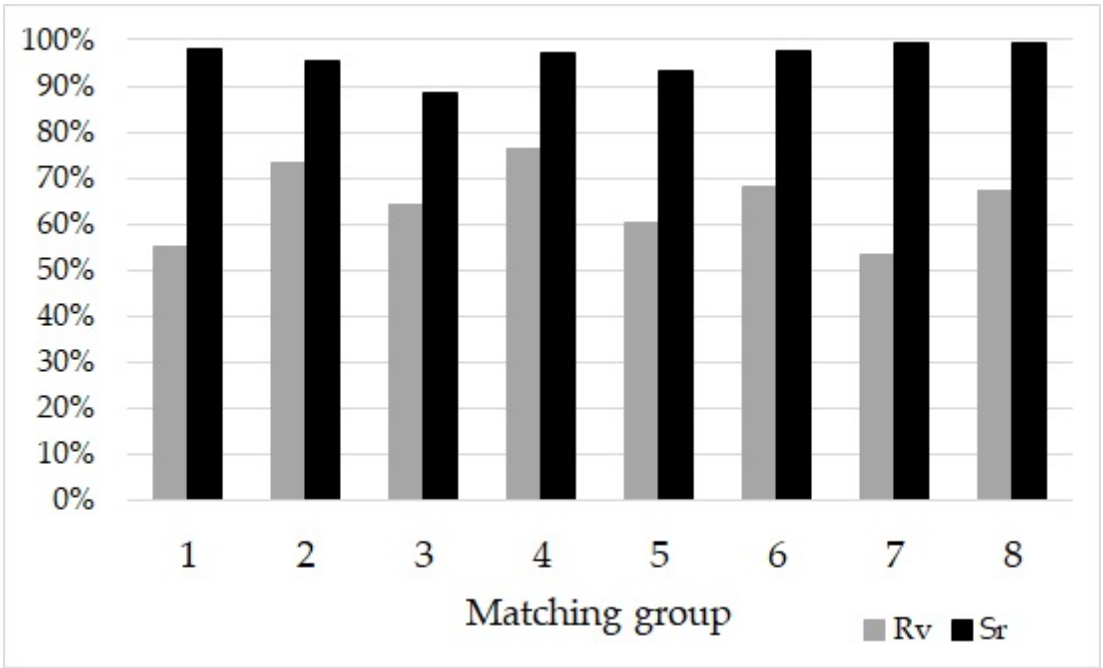


Figure 2.4: Frequency of revealing the higher card (Sr) or verifying the higher message (Rv) when a good hand is reported

As we conjectured, senders nearly always reveal the higher of the two cards in Sr (96.23% overall). In Rv, receivers check the higher of the two reports more often than the lower one in all matching groups (64.86% overall). However, the frequency of checking the higher report in Rv is significantly below the frequency of revealing the higher card in Sr in all matching groups[20]. Clearly, revealing the higher card is a more compelling strategy for senders than verifying the higher report is for receivers[21].

Result 5. *Senders nearly always reveal the higher of the two cards. Receivers usually check the higher of the two reports, but not quite as often.*

2.5.2.3 *Receivers' decision strategy*

When the sender reports a bad hand, the receiver nearly always rejects (over 96% of cases in both treatments). Focusing on cases where a good hand is reported, figure 2.5 presents the frequency with which receivers chose to accept for a given observed value. Specifically, the height of each dot depicts how often receivers accepted out of all rounds when they observed either a 1, 2,..., or 9. The lower graph presents the frequencies with which each value was observed in each treatment.

Comparing the Rv acceptance patterns with the Sr ones we see that the acceptance frequencies are greater in Rv than in Sr[22]. We can therefore say that receivers reject systematically more often when senders have verification control rather than when receivers have it. This is consistent with receivers in Sr holding more skeptical beliefs regarding the probability of a good hand than in Sr. Accordingly, we formulate the following result:

Result 6. *Receivers are more skeptical in the Sr treatment than in the Rv one.*

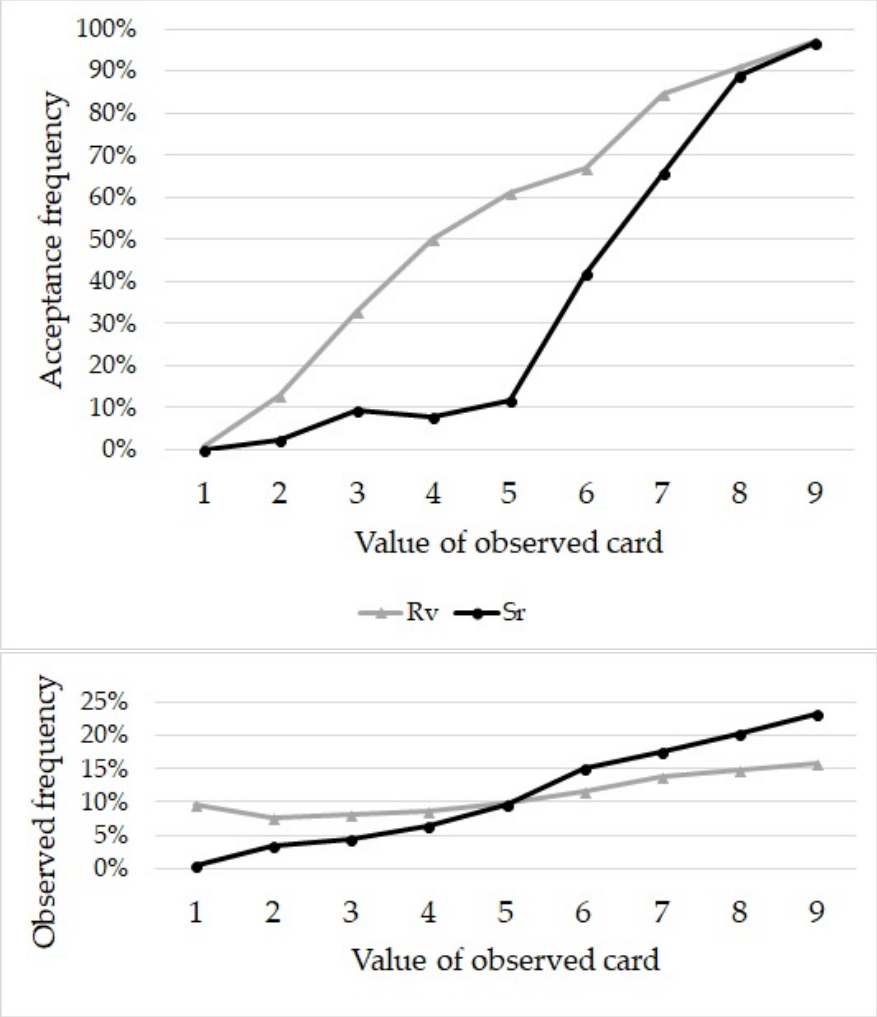
Even though the same random draws are used for both treatments, some differences in the distributions of the observed values (lower graph in figure 2.5) arise from different behavior across treatments. Specifically, due to differences in the revelation/verification strategy, the observed card in Sr is almost certain to be the highest of the two, while this is not the case in Rv.

20 The overall frequencies increase slightly (to 97.43% and 67.51% respectively) if we restrict ourselves to cases in which only one card or only one message is at least 7, but the general conclusion is unaffected. Those are the cases where the GR prediction is strict, since the receiver is predicted to accept cards with values of at least 7 and reject lower cards.

21 How often is the card with the highest actual value observed in Rv? There is a strong correlation between having the highest reported value and having the highest actual value. As a result, the probability of observing the highest card in Rv is also about 64%.

22 For some of the observed values the counts are too low to allow for meaningful testing (counts for low observed numbers in Sr are generally small). For each of the middle values 4-5-6-7, the difference is significant.

Consequently, the distribution of observed cards in Sr is very similar to the distribution of the maximum value out of the two cards, while the distribution in Rv is comparatively flat. The lower skepticism in Rv is thus partially justified: for a given value of the observed card, the hand is more likely to be good in Rv than in Sr.



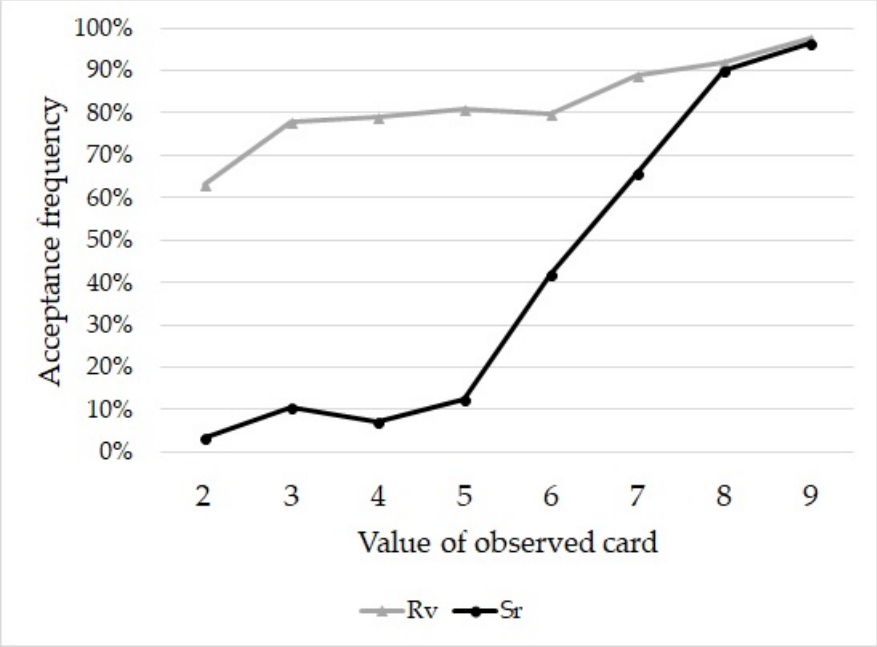
Note: 1245 counts for Sr; 1242 counts for Rv

Figure 2.5: Acceptance and observation frequencies conditional on the value of the verified card (reported good hands only)

Note however that the overall acceptance rates conditional on the sender claiming to have a good hand are very similar: 60% versus 62%. While the receiver is more likely to accept for a fixed value of the observed card in Rv, he is also more likely to observe lower cards in Rv (and lower cards are less likely to be accepted). Overall, the lower skepticism in Rv has no significant effect on the acceptance probability.

Figure 2.5 presents the acceptance strategy without taking the sender's message into account. The message makes little difference in Sr, since senders hardly ever misreport the card they reveal (less than 4% of reported good

hands involve misreporting the revealed card). It does make a difference in Rv, where in 30% of cases a misreport is observed. Receivers typically reject when observing a misreport (only 5% accept in Rv, 7% in Sr). This is optimal since only about 10% of such hands are actually good hands.



Note: 1199 counts for Sr; 862 counts for Rv

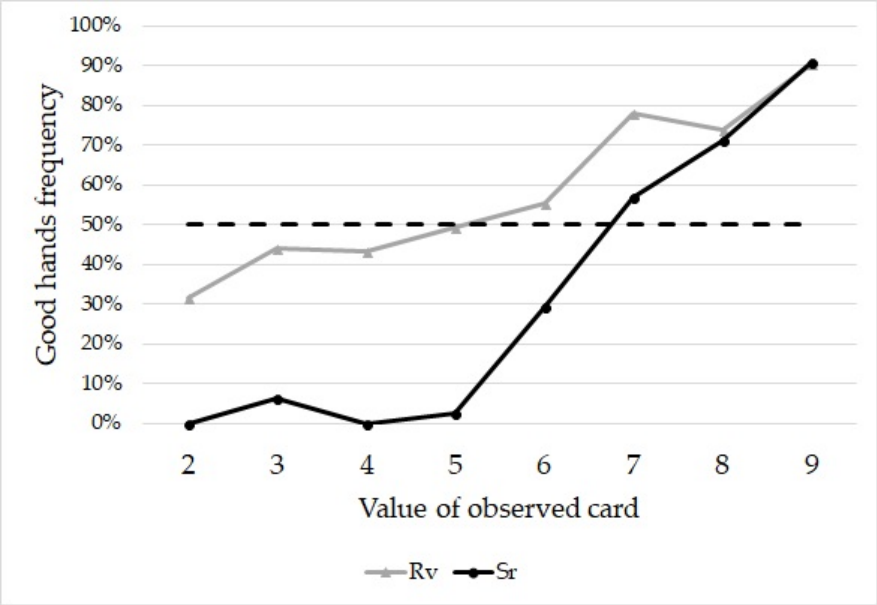
Figure 2.6: Acceptance frequencies when a good hand is reported and no misreport is observed

As figure 2.6 shows, the receiver is very likely to accept in Rv if he finds no reason to the contrary. This behavior is not as costly as one may think since a) some bad hands have been weeded out when a misreport is discovered (the sender's messaging behavior facilitates this) and b) the receiver's verification strategy implies that the card observed is not always the higher of the two. The actual frequency of good hands is shown in figure 2.7 (the figure includes a dotted line at 50% - above this line, it is optimal on average to accept, while below this line it is optimal on average to reject).

On average, in Sr it is optimal for the receiver to accept after observing a 7+, just as in the theoretical prediction (figure 2.7); however, the receiver's actual acceptance behavior is noisy around values 6-7 (see figure 2.6)[23]. In Rv however, given that the receiver sometimes checks the lower message and given

23 Figures 2.5 and 2.6 do not include information on the correlation between good hands and accepted hands for a given observed value. In principle, if the message sent by the sender is (unwittingly) informative (for example, by always inflating bad hands up to the bare minimum) it would be possible for the receiver to "see through the message" and use the message to improve the accuracy of the acceptance decision. The data suggest that this was not the case: conditional on a good hand being reported and holding the value observed constant, acceptance frequencies of good hands and bad hands are not significantly different. Hence, the noisy acceptance behavior corresponds, on average, to mistakes by the receiver.

that misreports are sometimes discovered and eliminated from the sample in the figure, the probability of a good hand remains close to 0.5 even for values as low as 3, such that the cost of an acceptance decision is small in terms of payoffs. Hence, even though the receiver is still too lenient for low observed values in Rv, this doesn't hurt the receiver too much (recall also that more misreports are detected in Rv, and the receiver is doing well for those hands, which are not included in figures 2.6 and 2.7).



Note: 1199 counts for Sr; 862 counts for Rv

Figure 2.7: Frequency of good hands given that a good hand was claimed and no misreport was observed

The figures give us some pointers on why average payoffs do not differ between treatments despite the differences in behavior.

Result 7. *(a) Receiver's acceptance strategy is noisy in Sr. In Rv, the receiver rejects too seldom when no misreport is discovered. However, this leniency is not very costly given other deviations by both the sender (reporting strategy, which makes it easier for the receiver to reject bad hands) and the receiver (verification strategy).*

(b) Even though receivers are more lenient in Rv conditional on the value observed, they are also more likely to observe lower values as a result of the verification strategy. As a result, overall acceptance frequencies and hence sender's payoffs do not differ between treatments.

We ran a probit regression for the acceptance decision and computed the marginal effects. We included only those observations where the reported values add up to at least 11 since receivers almost always reject when that is not the case (97.5% of bad hand reports are rejected in Rv and 96.4% in

Sr). The results are presented in table 2.6. First, we notice that increasing the value of the observed card has a significant positive effect on the acceptance probability in both treatments. Next, the value of the unverified report does not matter for receivers in Rv, but it has a significantly positive effect in Sr. The size of this effect is rather small but it is still surprising since messages should be completely ignored in Sr. Moreover, the receiver is more skeptical of claims that add up to 11 in Sr which again suggests that messages are not ignored in this treatment. In the next section we discuss whether this is an optimal strategy.

Table 2.6: Marginal effects from probit regression

	Dependent variable: Acceptance Probability	
	Rv	Sr
Value of observed card	0.111^{***}	0.221^{***}
	(0.025)	(0.023)
Value of unverified report	0.034	0.049^{***}
	(0.028)	(0.011)
No. of previous lies	-0.046^{**}	-0.034^{**}
	(0.020)	(0.014)
Discovered a lie in current period	-0.731^{***}	-0.336
(dummy)	(0.031)	(0.207)
Checked lower report	0.106	-0.216^{***}
(dummy)	(0.093)	(0.067)
Claimed sum = 11	-0.133	-0.112^{**}
(dummy)	(0.084)	(0.052)
Color (1=orange)	0.037	0.026
	(0.044)	(0.041)
Period	0.018	0.012
	(0.012)	(0.008)
Gender	-0.024	-0.009
(1=female)	(0.038)	(0.060)
Age	-0.015	-0.008
	(0.010)	(0.007)
Observations	1,242	1,245

Note: *p<0.1; **p<0.05; ***p<0.01
(Standard errors clustered at matching group level)

The number of times one has been previously lied to also has a small but significantly negative effect on the acceptance likelihood both in Sr and Rv. Discovering a lie in the current period, is a strong acceptance deterrent in Rv, but not in Sr. Admittedly, this almost never happens in Sr as senders do not reveal the cards they misreported. This, however, suggests that receivers in Rv

recognize that only those with a bad hand have an incentive to lie. Finally, we note that whether the low report was verified has no effect in Rv. In Sr though, revealing the low report has a strong and significantly negative effect on the acceptance probability. This is an indication that receivers in Sr understand that revealing the lower card is a strong signal of a bad hand.

2.5.3 *How can the receiver best protect himself?*

We answer this question by identifying the optimal strategy for the receiver given the sender's behavior for a fixed mechanism (Sr or Rv). We specify below which criteria we use for identifying these strategies for each mechanism.

2.5.3.1 *Receiver verifies*

Given the actual sender behavior in the Rv experiment, should the receiver check the card corresponding to the highest or to the lowest message? Moreover, should he pay attention to the unverified report? Lastly, which values should the receiver accept?

First, if the sender reports a bad hand the receiver can verify either report and should always reject as the report is almost certainly representative of the true state (in Rv, 1 out of 198 reported bad hands is good; in Sr, 3 out of 195 reported bad hands are good). Next, the receiver should check the report that yields the highest percentage of good hands conditional on the report being true. This is because, if the report is untrue, the hand will be a bad hand and the receiver will then reject it and make no errors. Here we assume that only the senders that have an incentive to lie (i.e. have a bad hand) actually do so[24]. If the report is true, the receiver minimizes the number of errors by maximizing the number of good hands he accepts. In other words, the higher the probability that the hand is good conditional on the report being true, the lower the number of errors that the receiver makes overall.

Equivalently, the receiver should check the report that is more likely to be false in which case it almost always represents a bad hand. This allows the receiver to take bad hands out of the sample by rejecting which increases the probability that the hand is good conditional on the message being true, and reduces the number of errors in case of acceptance (in cases where this probability is less than 50%, it is optimal to reject all hands and which message is checked is immaterial).

24 This is generally true as only 8% of senders lie about good hands in the Rv treatment (most likely due to error).

Table 2.7: Calculating the best response to the empirical distribution for Rv

Message	Absolute frequency of message	Absolute frequency of good hands given message	% Good hands given highest report is true	% Good hands given lowest report is true	Expected payoff from: check highest report & accept if true[25]	Expected payoff from: check lowest report & accept if true[26]	Expected payoff from: check either report & reject[27]
(9,9)	55	32	95%	95%	0.773	0.773	0.418
(9,8)	40	35	100%	97%	0.900	0.900	0.125
(9,7)	45	38	97%	95%	0.978	0.933	0.156
(9,6)	36	34	97%	97%	0.972	0.944	0.056
(9,5)	55	45	90%	94%	0.909	0.945	0.182
(9,4)	29	24	96%	86%	0.966	0.862	0.172
(9,3)	55	35	76%	80%	0.800	0.836	0.364
(9,2)	73	23	82%	36%	0.932	0.438	0.685
(8,8)	31	22	85%	85%	0.726	0.726	0.290
(8,7)	40	29	93%	87%	0.850	0.850	0.275
(8,6)	51	40	93%	88%	0.941	0.843	0.216
(8,5)	52	37	86%	80%	0.885	0.827	0.288
(8,4)	69	27	60%	55%	0.739	0.681	0.609
(8,3)	117	32	56%	35%	0.786	0.496	0.726
(7,7)	47	24	83%	83%	0.872	0.872	0.489
(7,6)	41	31	91%	94%	0.902	0.927	0.244
(7,5)	81	30	54%	55%	0.691	0.704	0.630
(7,4)	129	38	59%	38%	0.798	0.527	0.705
(6,6)	56	25	63%	63%	0.696	0.696	0.554
(6,5)	140	30	34%	36%	0.586	0.621	0.786
				Total expected payoff	0.804	0.714	0.492

25 Computed by averaging over the instances in which the highest report is untrue and the hand is bad and those where the highest report is true and the hand is good. For each such instance the receiver earns 1 point which is then weighted by the frequency of the corresponding message.

26 Computed by averaging over the instances in which the lowest report is untrue and the hand is bad and those where the lowest report is true and the hand is good. For each such instance the receiver earns 1 point which is then weighted by the frequency of the corresponding message.

27 Computed by adding the number of bad hands and then dividing by the frequency of the corresponding message.

Looking at table 2.7[28], it is clear that checking the highest report is a better choice than checking the lowest report for most messages since it leads to a higher percentage of good hands conditional on the report being true (in cases where it is not, the numbers are very similar). Admittedly, as long as this percentage is above 50%, the differences are small.

We also see that it is optimal for the receiver to reject all hands that send message (6,5): even if the observed value is as reported, the hand is still more likely to be a bad hand than not[29]. The only case in which the best response to the empirical distribution clearly differs from the GR equilibrium strategy is when the sender sends message (6,6). According to the strategy, these hands should be rejected. However, given the actual sender's strategy, the receiver should accept if the value of the observed card is indeed 6. Given the sender's actual behavior, such hands are more likely to be good than bad and it would be optimal to accept them[30].

If we pool all messages (excluding those where the two reported values are equal) we find that, conditional on a good hand being reported in Rv, the lower message is more likely to be true than the higher message in all matching groups. Hence, the receiver is better off by checking the higher message (the higher message is true in 67% of cases compared with 77% of the time for the lowest message). Conditional on the highest message being true, 73% of hands are good; conditional on the lowest message being true only 63% of hands are good.

The optimality of checking the highest report and accepting if it is true is observed also by directly comparing the expected payoff from this strategy with the other possible strategies (see last three columns of table 2.7). First, we note that apart from the case where the message is (6,5), the receiver earns significantly less by rejecting rather than accepting. Next, except for 4 cases (where the differences are very small), the receiver earns at least as much by checking the higher report than by checking the lower one and accepting if they are true. The expected payoff from the former strategy over all good hand messages is 0.804 which is significantly higher than 0.714 the overall expected payoff from the latter strategy ($p - value = 0.008$).

Result 8. *The GR verification strategy (i.e. checking the higher message) is optimal for the receiver given the sender's actual behavior in Rv.*

28 The table omits cases where the sender admits to having a bad hand (when it is optimal for the receiver to reject).

29 At the matching group level, this is true in 7 out of 8 matching groups.

30 If we look at each matching group, observing a 6 after a (6,6) corresponds to a good hand on average, but only in 6 out of the 8 matching groups.

Note that Result 8 does not make a claim about whether the GR acceptance strategy is optimal since it is better for the receiver to accept rather than reject when the message is (6,6).

Are messages about the unobserved claim informative in Rv?

The GR prediction for the Rv game states that a message is useful for the receiver only to the extent that it guides the verification strategy. Once a report has been checked, the value of the unverified report should play no role. From table 2.7, we observe that holding the value of the higher report constant, the probability of a good hand conditional on the value of the lower report does not follow an obvious trend. Hence, we conclude that the unobserved claim is uninformative in Rv. Nevertheless, we note that when the sum of the two reports is equal to 11 the incentives to check the higher report and accept if it is true are stronger than in any other case.

2.5.3.2 *Sender reveals*

Clearly, when the sender admits to having a bad hand, the receiver should reject. In cases where a good hand is reported, the optimal acceptance threshold for the receiver remains at 7. This isn't obvious because, if enough senders with bad hands tell the truth, it could be optimal to accept a 6 when the hand is claimed to be good. Table 2.8 shows that this is not the case: the presence of subjects who always (or sometimes) tell the truth does not affect the receiver's optimal acceptance threshold. This is because if the observed value is at least 7 the hand is more likely to be good, whereas if the observed value is at most 6, the hand is more likely to be bad.

Result 9. *The GR strategy is optimal for the receiver given sender's actual behavior in Sr.*

Are messages informative in Sr?

The receiver must ignore the message in equilibrium in Sr. However, it could be possible that the receiver can protect his interests better by taking the message into account, given the actual sender behavior in the experiment. When the sender admits to having a bad hand, the receiver should act accordingly and reject. However, as mentioned earlier (see figure 2.3), in most of these cases the receiver would have been able to make the right decision by conditioning only on the value of the card displayed.

What about cases where the displayed card + the claim about the value of the other card add up to at least 11? Can the receiver do even better by taking the exact message into account? It is possible that a sender who has a bad

Table 2.8: Calculating the best response to the empirical distribution for Sr

Observed value	% Good hands
1	0%
2	2%
3	6%
4	0%
5	3%
6	29%
7	57%
8	71%
9	91%

Note: Table considers only the cases where the sender claims to have a good hand.

hand tends to inflate the lower card up to a level where the two cards add up to exactly 11. The receiver then may be justified in mistrusting reports that add up to exactly 11. Perhaps anticipating this possibility, the sender may be disproportionately likely to inflate the lower card until the two cards add up to exactly 11+1.

We may first ask whether, holding the value of the observed card constant, a message where the two cards add up to exactly 11 or 12 is bad news for the receiver. There is some evidence to this effect[31]. For observed values of 8 and 9, even though seeing reports that add up to exactly 11 or 12 reduces the probability of a good hand, it is still more likely to be a good hand than not, hence the optimal receiver strategy is not affected. For observed values of 7, the receiver would be better off on average by rejecting when seeing reports that add up to 11 or 12, although the difference in receiver payoffs between accepting and rejecting is small and statistically insignificant.

Overall, there is some evidence that, holding the observed value of the card constant, hearing that the sum of the reports barely makes the threshold is bad news for the receiver in that it reduces the probability that the hand is good. However, this has little or no effect on the receiver's best response.

31 Conditional on the observed card being 7, 8 or 9, the proportions of good hands are 57%, 71% and 91% respectively. These proportions drop to 49%, 54% and 78% when the claimed values add up to 11. When the message adds up to 12, these proportions drop to 43%, 52% and 90% respectively. These differences are statistically significant at the 5% level (two-sided Wilcoxon signed ranks tests, separately for each of the three observed values and comparing reports that add up to either 11 or 12 with reports that add up to 13 or more).

2.5.4 *Comparison of payoffs with the GR equilibrium*

Lastly, we ask how costly the observed deviations from the GR prediction are. Average payoffs are not significantly different from the GR equilibrium prediction except for one case: the receiver's payoff in Sr which is significantly below the prediction. Tables 2.9 and 2.10 present the payoff comparisons. The GR average payoffs are computed given the actual realization of good and bad hands while assuming that both players behave according to the optimal strategies.

Table 2.9: Average payoffs: Rv vs. GR equilibrium predictions

	Rv	GR	p-value
Sender	0.541	0.554	0.328
Receiver	0.782	0.813	0.187

Table 2.10: Average payoffs: Sr vs. GR equilibrium predictions

	Sr	GR	p-value
Sender	0.525	0.554	0.203
Receiver	0.767	0.813	0.008

Result 10. *The receiver's payoff in Sr is significantly below the equilibrium prediction*[32].

It may seem surprising that the receiver's payoff is significantly below GR in the Sr case, while the difference between Sr and Rv (Result 2) or between Rv and GR is not significant. This is because, at the matching group level, receiver's payoffs are always below GR in Sr, while there is no consistent ranking in the comparison of Rv with either Sr or GR[33].

2.6 A DIFFERENT PARAMETRIZATION

We also test behavior in the same two games but with a different threshold (T) defining good hands. In this implementation of the Sr and Rv games the two cards have to add up to at least 9 in order to represent a good hand. We implemented these treatments as a robustness check. One can also regard this

32 This is a very robust result; it still holds when selecting only the last 25, 20, 15 or 10 rounds. The receiver's payoff is consistently below the theoretical prediction in all matching groups (except for a tie in one matching group in the last 10 rounds).

33 See Appendix A.2 for the observed and predicted average payoffs per matching group.

as a test for whether different levels of conflict lead to different results since in the case where T=9 there is relatively less conflict between the sender and the receiver than when T=11. This is because in the former case there are more good hands than bad hands whereas in the latter case the opposite is true. The level of conflict of interests is an important element in cheap talk games as the lower the conflict the higher the amount of information transmitted in equilibrium[34]. In our games, this plays no role as sender's preferences are state independent. Consequently, the GR prediction is the same except for the fact that the receiver will insist on seeing a 6 rather than a 7 in order to accept.

We used the same experimental design and procedures. There were 180 subjects in total (66% of subjects in Rv were female while in Sr this percentage was 65%). One session in the Rv treatment had 20 instead of 24 subjects. The corresponding session in Sr as well as one other also had only 20 subjects. In the following analysis, for the sessions in which the Rv treatment had more subjects than the Sr one, we drop the equivalent extra observations in Rv to maintain comparability between the underlying draws. The main analysis is therefore performed on 8 independent observations, but now 4 of them comprise of 10 individuals rather than 12.

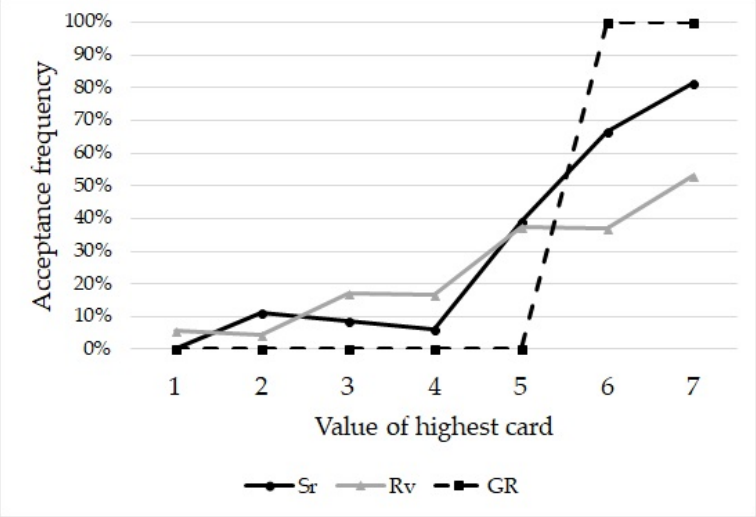
In this section, we present the main results regarding our two hypotheses using the data from these additional treatments. Recall that our two main hypotheses are that outcomes are closer to GR predictions in Sr than in Rv and that consequently, payoffs are higher in the former than in the latter. We begin the analysis of the T=9 treatments, by looking at the receivers' acceptance strategy (figure 2.8 and table 2.11) to identify how close each treatment outcomes are to GR predictions. We note a similar pattern as in T=11: receivers' acceptance strategy is closer to GR predictions in Sr than in Rv since they reject more low-valued hands and accept more high valued ones.

The difference is less strong than in T=11 though and this is reflected also by the hit rate analysis: 84.01% of outcomes are predicted by the GR equilibrium in Rv and 85.76% in Sr. This difference is not statistically significant ($p - value = 0.172$). Both players are better off in Rv than in Sr but this difference is not statistically significant (table 2.12) at the 5% level. It is though worth noting that this time, the difference in receiver's average payoff is significant at a 10% level. This is an indication that the Rv environment can make receivers better off conditional on the conflict level being low enough.

With regards to specific strategies, the results from the T=9 treatments support the robustness of the T=11 findings. First, we note that senders' reporting strategies are in line with Result 4. For bad hands that would be accepted in equilibrium, senders in Sr report a good hand while keeping the highest of

34 When interests are perfectly aligned, communication in equilibrium is perfectly informative (Crawford and Sobel, 1982).

(a) Bad hands



(b) Good hands

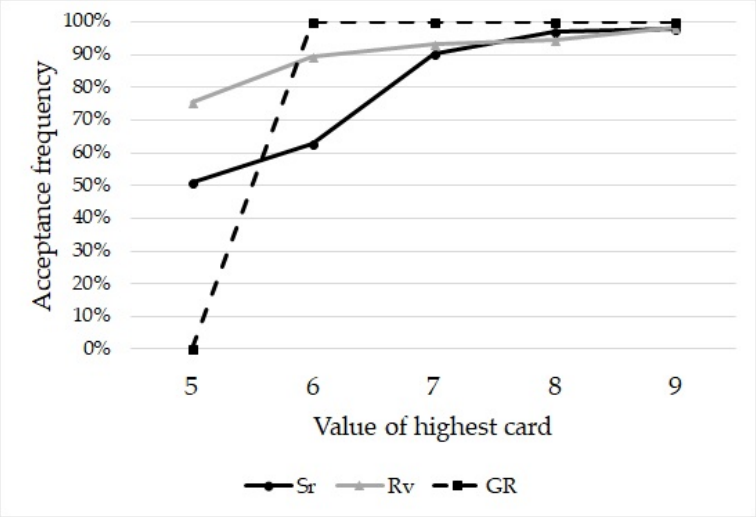


Figure 2.8: Acceptance frequencies for a given value of the highest card (T=9)

Table 2.11: Acceptance frequencies for pooled data (within 6 categories)

Type of hand		Freq. GR	Freq. Sr	Freq. Rv	count	p-value
(a) Bad hands	No high card (1-5)	0.000	0.173	0.211	365	0.156
	One high card (6-9)	1.000	0.721	0.430	86	0.031
	All bad hands	0.191	0.277	0.253	451	0.438
(b) Good hands	No high card (1-5)	0.000	0.509	0.754	57	0.031
	One high card (6-9)	1.000	0.911	0.948	812	0.234
	All good hands	0.934	0.885	0.936	869	0.055

Note: Counts are for each treatment.

the two reports truthful in 85% of the cases, while in Rv this happens only in 53% of cases which is a significantly lower percentage ($p - value = 0.016$). Moreover, consistent with Result 5, senders reveal the higher card system-

Table 2.12: Average payoffs in Sr and Rv when T=9

	Sr	Rv	p-value
Sender	0.678	0.701	0.266
Receiver	0.810	0.872	0.086

atically more often in Sr (98.56%) than receivers verify the higher report in Rv (82.64%). The difference in the corresponding frequencies is statistically significant ($p - value < 0.001$). Next, for any given verified/revealed value receivers are more skeptical in Sr than in Rv, in line with Result 6. However, the overall acceptance frequencies (77.59% for Rv and 75.38% for Sr) do not differ significantly in the T=9 treatments either.

Finally, the GR strategy is overall optimal for the receiver given the senders' behavior in the treatments where T=9. In the Sr case, the probability that the hand is good given that the observed value is at least 6 is always higher than 50%. If the observed value is at most 5, this probability is always less than 50% which makes accepting hands where the revealed card is at least 6 optimal.

In the Rv case, recall that the optimal strategy depends on the probability of a good hand conditional on the verified report being true (the higher the better) and on the probability that a given report is true (the lower the better). Focusing on reported good hands, the probability of a good hand conditional on the highest report being true is always at least as high as when conditioning on the lowest report being true. Moreover, the highest message is true in 68% of the cases whereas the lowest one in 79%. Consequently, following the strategy of checking the highest report when a good hand is reported and accepting if it is true gives a higher expected payoff than checking the lowest report and accepting when it is true. This holds for all good hand messages including the case of (5,5), for which GR predicts that the hand should be rejected[35]. In this case, like in T=11 when the message was (6,6), the empirical distribution reveals that the probability of a good hand conditional on the verified report being true is higher than 50%. This means the receiver is better off accepting these hands rather than rejecting them.

In summary, when T=9 we find a pattern of results very similar to the case where T=11. Senders' behavior is closer to GR predictions in Sr rather than Rv while receivers are less skeptical in Rv than in Sr for a given observed value. This, however, does not decrease players' average payoffs because in Rv senders deviate in ways that reward a more lenient response from receivers. Nonetheless, the GR strategy is optimal under these parameters as well, with one exception: the Rv case where the message is (5,5).

35 See Appendix A.3 for the data regarding the receiver's optimal action given the empirical distribution in T=9.

2.7 CONCLUSION

We have studied games of strategic information transmission between a sender and a receiver in which cheap talk claims are partially backed by evidence. Such games reflect many naturalistic environments like buyer-seller interactions where the buyer is exposed to advertisements or listens to a selling pitch but can also test the products (e.g. downloading samples of software or test-driving cars). Other situations involve policy makers that are pressured by self-interested lobbyists to implement certain legislation but who can also acquire (limited) facts regarding the lobbyists' claims. These situations have received little attention in the experimental literature so far.

We focused on two settings based on two theoretical models of Glazer and Rubinstein (2004, 2006). These differ in the amount of control the receiver has over the information space. In one setting (Rv) the receiver can choose which cheap talk claim to verify, while in the other (Sr) the sender has this choice. Although theoretically the two environments have the same optimal prediction for the receiver, in Sr this prediction is more compelling than in Rv. We do not observe significant differences in terms of average payoffs between treatments. Nonetheless, behavior differs significantly across games. Specifically, the sender's message strategy is further from optimal predictions in Rv than in Sr. This allows the receiver to deviate from his optimal strategy in ways that are profitable: by rejecting more bad types and accepting more good ones. Overall, the receiver is less skeptical in the Rv treatment than in Sr and this pays off with respect to the equilibrium benchmark. This is a surprising result as a priori there is no reason for receivers to be less skeptical in Rv than in Sr. It is unclear whether they are knowingly taking advantage of senders' deviations or whether the higher leniency is an effect of the increased control over their information space. Some recent evidence tentatively points to the latter: Penczynski and Zhang (2016) find that when buyers can choose the seller whom to buy from (competition setting) they act less skeptical than when they do not have this choice (monopoly setting).

With respect to which environment is better for the receiver, there is no clear winner. While the deviations from the optimal GR strategy are advantageous for the receiver in Rv, the Sr environment promotes behavior closer to the GR prediction which can also benefit the receiver. From this perspective, an interesting investigation could be to control for receivers' level of sophistication and see whether less sophisticated players do better in Rv where errors are less costly than in Sr. Finally, we looked at which strategy is optimal for the receiver given the senders' behavior in the experiment. Our results support the theoretical prediction that following the GR strategy is the receiver's best option in both the Sr and the Rv games.

BIBLIOGRAPHY

Abeler, J., Nosenzo, D., and Raymond, C. (2016). Preferences for truth-telling. *CESifo Working Paper Series No. 6087*.

Alonso, R. and Câmara, O. (2016). Persuading voters. *American Economic Review*, 106(11):3590–3605.

Blume, A., Lai, E. K., and Lim, W. (2017). Strategic information transmission: A survey of experiments and theoretical foundations. *Mimeo*.

Cabrales, A., Feri, F., Gottardi, P., and Melendez-Jimenez, M. (2016). Can there be a market for cheap-talk information? Some experimental evidence. *Mimeo*.

Carroll, G. and Egorov, G. (2017). Strategic communication with minimal verification. *Mimeo*.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Forges, F. and Koessler, F. (2005). Communication equilibria with partially verifiable types. *Journal of Mathematical Economics*, 41(7):793–811.

Forges, F. and Koessler, F. (2008). Long persuasion games. *Journal of Economic Theory*, 143(1):1–35.

Forsythe, R., Isaac, R. M., and Palfrey, T. R. (1989). Theories and tests of "blind bidding" in sealed-bid auctions. *The Rand Journal of Economics*, pages 214–238.

Glazer, J. and Rubinstein, A. (2004). On optimal rules of persuasion. *Econometrica*, 72(6):1715–1736.

Glazer, J. and Rubinstein, A. (2006). A study in the pragmatics of persuasion. *Theoretical Economics*, 1:395–410.

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

Hagenbach, J. and Perez-Richet, E. (2018). Communication with evidence in the lab. *CEPR Discussion Paper No. DP12927*.

Hart, S., Kremer, I., and Perry, M. (2017). Evidence games: Truth and commitment. *American Economic Review*, 107(3):690–713.

Jin, G. Z., Luca, M., and Martin, D. (2015). Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure. *NBER Working Paper No. 21099*.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

King, R. R. and Wallin, D. E. (1991). Market-induced information disclosures: An experimental markets investigation. *Contemporary Accounting Research*, 8(1):170–197.

Li, Y. X. and Schipper, B. C. (2018). Strategic reasoning in persuasion games: An experiment. *Mimeo*.

Lipman, B. L. and Seppi, D. J. (1995). Robust inference in communication games with partial provability. *Journal of Economic Theory*, 66(2):370–405.

Lipnowski, E. and Ravid, D. (2017). Cheap talk with fixed agenda. *Mimeo*.

Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821.

Minozzi, W. and Woon, J. (2013). Lying aversion, lobbying, and context in a strategic communication experiment. *Journal of Theoretical Politics*, 25(3):309–337.

Penczynski, S. P. and Zhang, S. (2016). Disclosure of verifiable information under competition: An experimental study. *Mimeo*.

Serra-Garcia, M., van Damme, E., and Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, 73(1):244–261.

3

CHEAP TALK AND EVIDENCE IN THE LAB

3.1 INTRODUCTION

In many real sender-receiver types of environments, the receiver typically evaluates two sources of information. On the one hand he is presented with cheap talk, i.e. non-verifiable information; on the other hand, he receives evidence which more often than not is partially informative. One such example is the FDA drug approval process. In this case, the cheap talk consists of the lobbying activities that the pharmaceutical companies perform towards the FDA. The hard evidence arises when, in order to get approval, the company runs tests to showcase the safety of the drug. The structure of this testing procedure (which tests to include and how to run them) is decided by the applicant company. Due to resource (time and money) constraints, these tests cannot be fully revelatory about the properties of the drug. Hence, one can characterize this piece of information the FDA needs to assess as partially informative evidence.

How do individuals interpret hard evidence when this is accompanied by cheap talk? A plethora of empirical studies have analyzed the effect of communication in experimental games. The typical result is that communication improves outcomes and at the very least, has no effect in the games it has been analyzed. Positive effects are found in the dictator and ultimatum games (Mohlin and Johannesson, 2008; Roth, 1995; Zultan, 2012) while in trust games it increases trustworthiness (Charness and Dufwenberg, 2006, 2011). In sender-receiver games with asymmetric information, the effect of communication is not so clear. A relatively recent finding among such studies is that communication can increase receivers' skepticism, making them refrain from an action they would have otherwise taken (Clots-Figueras et al., 2015; Ert et al., 2014). The effect of cheap talk on the interpretation of hard evidence has not yet been studied though. In absence of communication when full but voluntary disclosure is possible, receivers seem to be insufficiently skeptical about undisclosed information (Jin et al., 2015), whereas when voluntary partial disclosure is possible they seem to correctly assess the evidence presented (Hagenbach and Perez-Richet, 2018).

In this study, we look at whether cheap talk has an effect on decision makers' ability to make good decisions in environments with mandatory partial

disclosure. To this end, we employ a lab experiment using a 4-stage sender-receiver game based on the one discussed by Glazer and Rubinstein (2006). In the first stage of the game, the sender is dealt a 2-card hand, one card orange, one card blue. Each card can take a value between 1 and 9; each combination of values is equally probable. If the cards add up to at least 9, the hand is good; otherwise, the hand is bad. In the second stage, the sender sends a strict-form message to the receiver reporting the values on the two cards. The reported values need not be the true ones, but are bound between 1 and 9. Next, he chooses one of the two cards to reveal to the receiver. The receiver observes the message and the revealed card and decides whether to accept or reject.

Standard theory predicts that cheap talk should not affect the equilibrium payoff set. In the receiver's optimal equilibrium (the GR equilibrium) the receiver ignores the message and conditions his acceptance strategy only on the revealed card. This is because the sender's equilibrium strategy is to always reveal the higher card (and randomize in case of equality). Consequently, the receiver maximizes his payoff if he always rejects when the revealed card is less than 6, and always accepts when the revealed card is at least as high as 6.

We implement a low-conflict environment (the cards have to add up to at least 9) and a high-conflict one (the cards have to add up to at least 11). We find weak evidence that the receiver's payoff is lower in the low-conflict case when cheap talk is present. No significant difference is detected in the high-conflict case. The main driver of this negative effect of communication is skepticism when observing high-valued evidence. In a follow-up treatment we find that this effect is partially mitigated by giving the receivers the option to see the message or not. We attribute the observed reduced skepticism to the effect of increased control over one's information space which increases receivers' confidence that chances of being deceived are lower.

The rest of the chapter is structured as follows: Section 3.2 gives a detailed overview of the relevant literature, Section 3.3 presents the experimental design and procedures; this is followed by the results analysis in Section 3.4. Section 3.5 introduces the follow-up treatments and the relevant results while Section 3.6 concludes.

3.2 RELATED LITERATURE

This study is related to two strands of literature: cheap talk and disclosure of information (disclosed information represents what has been previously referred to as evidence). On the one hand, we depart from the cheap talk literature by introducing partially-verifiable messages. On the other hand, we

differ from the voluntary disclosure research through including non-verifiable claims in sender's action space. In addition, the type of disclosure in our experiment is not voluntary, but has a compulsory and specific format. This section will discuss the relevant existing research in both of these areas.

3.2.1 *Cheap talk*

The seminal paper of Crawford and Sobel (1982) has offered a benchmark for the theoretical analysis of strategic information transmission. In their model, the sender (informed party) sends a message to the receiver (uninformed party) about the state of the world after which the receiver takes a payoff relevant action. The message has no direct effect on payoffs; the message space is equal to the state space but the content of a message is not bounded to the realized state. Both players' preferences regarding the receiver's action depend on the state but differ in that the receiver always wants to match the state while the sender always prefers a weakly higher action. The size of this preference difference represents the size of the conflict of interests between the two parties. This game has multiple equilibria. However, in all of them the message strategy allows the receiver to partition the state space in a way such that he can take the same action for all elements of a partition. Moreover, the authors show that as the conflict of interests increases, the information transmitted decreases, i.e. the receiver is able to make fewer and fewer partitions.

Cai and Wang (2006) put this theory to a test using a lab experiment. They find support for the prediction that the higher the conflict of interests the less information is communicated. However, their results show that when compared with the most informative equilibrium predictions[1], senders tell the truth more often while receivers follow the messages more frequently than theory predicts. Average payoffs are close to the most informative equilibrium prediction though.

Clots-Figueras et al. (2015) test the effect of communication in a modified investment game. In this game, the investor decides how much of his endowment to invest. This amount is multiplied by a positive parameter after which, the investee decides how much to return to the investor. The authors introduce asymmetric information about the value of the investment multiplier, communicating this value only to the investee. Before the investor takes an action, the investee sends a message to the investor about the value of the multiplier. According to the main result, messages make investors more pessimistic (they invest less) when the distribution is known and more optimistic

1 The most informative equilibrium is the one in which the maximum amount of information is transmitted. This amount differs depending on the size of the conflict of interests.

(they invest more) when it is unknown. There is no reported effect of communication on average payoffs.

The type of messages used in the above studies can be labeled as strict-form messages as senders' communication is bound to a particular space (typically the state space). Another strand of literature focuses on free-form communication. This incorporates the situations where the message space is undefined. Consequently, the content of senders' messages can be anything from numbers to sentences unrelated to the game.

In the free-form cheap talk literature, communication is generally found to lead to positive effects within the games it has been studied. It increases donations in dictator games (Mohlin and Johannesson, 2008) and agreement frequencies in ultimatum games (Roth, 1995; Zultan, 2012). In trust games, communication increases the proportion of trustors who choose to trust (Charness, 2000; Charness and Dufwenberg, 2011). In sender-receiver games, cheap talk leads to efficiency gains, but these come at the expense of the receiver (Forsythe et al., 1999).

Ert et al. (2014) implement an experiment similar to ours in some respects. Their study is based on the "hidden card" game (Rubinstein, 2012) where the state of the world is represented by two cards drawn randomly by the seller from a deck of cards with values from 1 to 100. The buyer is then informed about the value of the lower card and has to decide whether to buy a product or not. Like us, the authors compare treatments with communication and without. In the communication treatment, before the buyer makes a decision, the seller sends a free-form message to the buyer in order to convince the buyer to purchase the good. Since the seller does not control the card that is revealed, there are no strategic considerations regarding the value of the unobserved card and the buyer can always use Bayes' rule to estimate it. Moreover, messages should have no effect on the buyer's decision. Their results show that cheap communication leads to a decrease in buyers' propensity to buy the good from the seller. They attribute this result to the fact that communication puts buyers in the role of a persuadee, which highlights the conflict of interest and reduces inter-personal trust. We differ from this study by using two uncorrelated state dimensions (cards) and giving the seller control over which card is disclosed. Moreover, we use strict-form communication. This is helpful for a clearer analysis of the strategic use and effect of different types of messages.

3.2.2 *Disclosure*

The game we use was first discussed by Glazer and Rubinstein (2006). In their setting the sender cannot misreport the values of the two state-dimensions which translates into a game of partial mandatory disclosure, where cheap talk is not present. Nevertheless, as discussed in chapter 2, the receiver's best equilibrium has the same outcome irrespective of whether communication is present or not: hands with a 6+ (7+) will always be accepted.

Most of the disclosure literature looks at voluntary disclosure settings. These are sender-receiver games similar to the Crawford and Sobel (1982) cheap talk game in terms of the information structure and players' preferences. The difference is that now, the sender cannot send untruthful messages. His action is either to disclose his true type or to stay silent. According to (Milgrom and Roberts, 1986) these games are characterized by all sender types choosing to disclose due to receivers assuming the worst about non-disclosure (the unravelling principle). This concept does not apply to our game where senders do not have a choice of disclosing or not. Nonetheless, receivers' inferences about undisclosed information in these games can be informative as to how close to theoretical predictions they would behave in our experiment.

Jin et al. (2015) test the unravelling principle (Milgrom and Roberts, 1986) in a lab experiment using a standard one-shot voluntary disclosure game between a receiver and a sender. The authors find that players' behavior is not in line with equilibrium predictions: receivers do not infer the worst about undisclosed information and consequently, senders disclose too seldom. Feedback helps receivers correct their beliefs about non-disclosure and thus, senders' behavior converges to full disclosure over time.

Penczynski and Zhang (2016) investigate the effect of competition on voluntary disclosure in a buyer-seller environment. They find that introducing competition between sellers increases the amount of information revealed. At the same time, giving buyers the option to choose the seller they would trade with reduces buyers' skepticism. This result is in line with the findings from our follow-up treatments. Chapter 2 also documents a skepticism-reducing effect of increased information control. Specifically, in that experiment we find that giving receivers the option to choose which reported dimension to check reduces skepticism about observed evidence.

Hagenbach and Perez-Richet (2018) also implement a lab experiment using a sender-receiver game with partial disclosure and bi-dimensional states. The main difference with our setting is that in their experiment disclosure is voluntary. Senders, can choose to disclose 0, 1 or both of these cards. No cheap talk messages are present either. Their results show that receivers make correct

inferences about the observed disclosure actions and act upon these in a consistent manner. This reconciles our findings in the treatment where senders are not allowed to send messages and only mandatory partial disclosure is permitted.

3.3 EXPERIMENTAL DESIGN

In this section we begin by describing the structure of the benchmark game we implemented in the experiment. Then, we present how this game was adapted to test our hypotheses and the experimental procedures.

3.3.1 *The game*

The interaction takes place between two players: a sender and a receiver. The sender is informed about the state of the world, while the receiver is not. The timing of the game is as follows:

t1. The sender is dealt a 2-card hand: one card orange (o) and one card blue (b): $(o, b)\epsilon\{1, ..., 9\} \times \{1, ..., 9\}$. A total $(o + b)$ of $T\epsilon\{9, 11\}$ or higher represents a good hand, otherwise the hand is bad. The value of T is common knowledge from the start of the game.

t2. The sender reports the values of the cards (possibly untruthfully) to the receiver through sending a bi-dimensional message. The message space is equal to the state space.

t3. The sender chooses one of the two cards to reveal to the receiver.

t4. The receiver observes the message and the true value of the revealed card and decides whether to accept or reject.

Payoffs:

Table 3.1: Payoff Matrix (sender's payoff listed first in each cell)

	Receiver accepts	Receiver rejects
Sender has a good hand	(1, 1)	(0, 0)
Sender has a bad hand	(1, 0)	(0, 1)

From table 3.1 we note that incentives are such that the sender would always like the receiver to accept, whereas the receiver would like to accept only good hands and reject only bad ones.

Sequential equilibria:

As described in chapter 2, the above game has multiple sequential equilibria[2]. The one that stands out is the equilibrium in which messages are ignored and the sender always reveals the higher of the two cards. The receiver then accepts if the observed value is at least 6 when T=9 and at least 7 when T=11. Glazer and Rubinstein (2006) prove that when messages are not allowed, this equilibrium leads to the receiver's best equilibrium payoff. Adding messages does not change this. Hence, we will refer to this equilibrium as the GR equilibrium.

Moreover, when messages are not allowed, this equilibrium survives iterated elimination of weakly dominated strategies (see Appendix B.1). Admittedly, solving the game through this method is a lengthy process, requiring more than 3 rounds of iterated dominance in both T=9 and T=11 cases. Costa-Gomes et al. (2001) show that when this is the case, compliance with equilibrium predictions is not particularly high. Therefore, we cannot claim that the game with messages which is not dominance solvable, will lead to behavior closer to the GR predictions. We therefore hypothesize that communication should have no effect for either conflict level. We formally state the hypotheses after describing the experimental design.

3.3.2 *Treatments and hypotheses*

We implement a 2x2 experimental design to test the effect of communication. Our main variable is the presence of messages. We also vary the value of the threshold defining if a hand is good or bad. The 4 treatments are summarized in table 3.2.

Table 3.2: Main treatments

Tr. name	Messages allowed?	Threshold value (T)
CM_9	YES	9
No_CM_9	NO	9
CM_11	YES	11
No_CM_11	NO	11

A threshold value equal to 9 represents an environment where the level of conflict is relatively low as there are more good hands than bad hands. Put differently, when T=9, there are more states in which the interests of the sender and those of the receiver coincide. When T=11, conflict is relatively high as there are more states in which the interests of the two players are in

2 In chapter 2, this game is referred to as the *Sender reveals* game.

opposition. We implement two different levels of conflict to check whether the finding from the cheap talk literature that the lower the conflict of interests the higher the amount of informative messages (Cai and Wang, 2006) is relevant also when evidence is present. Since in our game unverified messages are ignored in equilibrium, we should not observe any difference in the amount of information transmitted conditional on the hand being bad.

Given this experimental design coupled with the theoretical predictions, we hypothesize the following:

Hypothesis 1 (behavior): The sender's revelation strategy and receiver's acceptance behavior are not different between the No_CM_ and the CM_ treatments irrespective of the value of T.

Hypothesis 2 (outcomes): The players' average payoffs in the treatments where messages are not allowed (No_CM_9 and No_CM_11) are not different than their average payoffs in the corresponding treatments with messages (CM_9 and CM_11).

3.3.3 *Procedures*

The experiment was conducted at the CeDEx laboratory within the University of Nottingham (UoN) using the UoN undergraduate and graduate student pool. We ran 4 sessions for each of the 4 treatments. Participants were recruited using ORSEE (Greiner, 2004) and assigned to one session at random. There were 372 participants: 88 in CM_9, 92 in No_CM_9, 96 in CM_11 and 96 in No_CM_11. In each session, participants played the game for 30 rounds, being randomly re-matched after every round. The re-matching was done within a matching group (MG) made up of half the participants in that session. This gives us 8 independent observations per treatment. There were 24 players in each session with the following 3 exceptions which had 20 participants: the first and last session of CM_9 and the last session of No_CM_9. To maintain comparability, in the following analysis of the data we drop the extra observations in the corresponding treatments for the T=9 case. We explain this in more detail in Section 3.4. The gender distribution was similar across treatments: 65% of subjects were female in CM_9, 69% in No_CM_9, 65% in CM_11 and 50% in No_CM_11.

After all participants within a session were seated, the instructions were read out loud by an experimenter[3]. Afterwards, players were randomly as-

3 See Appendix B.2 for full instruction of the No_CM_9 treatment. The instructions of the CM_11 treatment are presented in Appendix A.4.2. Their different T counterparts are identical except for the description of what constitutes a good hand.

signed to one of the two player roles. In the instructions, the two roles were titled in an abstract manner: the sender was named "Person A" and the receiver "Person B". Participants interacted through a computerized interface which was programmed in z-Tree (Fischbacher, 2007). As described in table 3.2, the difference between the CM_ and No_CM_ treatments consists in whether messages were present. The format of the messages that were sent in the CM_ treatments is depicted in figure 3.1; senders had to fill each blank box with a number between 1 and 9.



Figure 3.1: Message format

The distribution of the realized cards was identical between the CM_ and No_CM_ treatments. In order to do so, we randomly generated this distribution before the start of the sessions. In the T=9 treatments there were 869 good hands and 451 bad ones, while in the T=11 treatments we had 631 good hands and 809 bad ones. At the end of each round, participants received full feedback regarding the realized card values and the decisions each player made.

3.4 RESULTS

In this section we present the results of the experimental data analysis. We begin with players' behavior followed by how it translates to average payoffs. We do not observe any learning effect (behavior in the first half of the periods is similar to that in the second half) and so we treat each period as equally relevant for the analysis. The data for the CM_ treatments is presented in chapter 2 as well[4]. Nonetheless, for completeness, it is necessary to reiterate some of those findings. Unless otherwise specified, all reported p-values correspond to Wilcoxon signed-rank tests clustered at the matching group level [5].

4 Since the treatments were run simultaneously, this did not affect the random assignment of participants across treatments.

5 As the data in the two treatments with T=9 was not balanced, to maintain comparability of CM_9 and No_CM_9 we dropped the observations in No_CM_9 corresponding to the 2 missing sender-receiver pairs from the first session of CM_9. In the case of the last session of CM_9 where 2 pairs were missing as well, the corresponding No_CM_9 session was run with 20 instead of 24 subjects so we did not have to drop observations ex-post. This was necessary as the addition of more good hands or more bad hands would change the realized level of conflict of interests which would make the treatments non-comparable.

3.4.1 *Reporting strategy in CM_ treatments*

Figure 3.2 presents the frequency of each combination of realized (left panels) and reported (right panels) values.

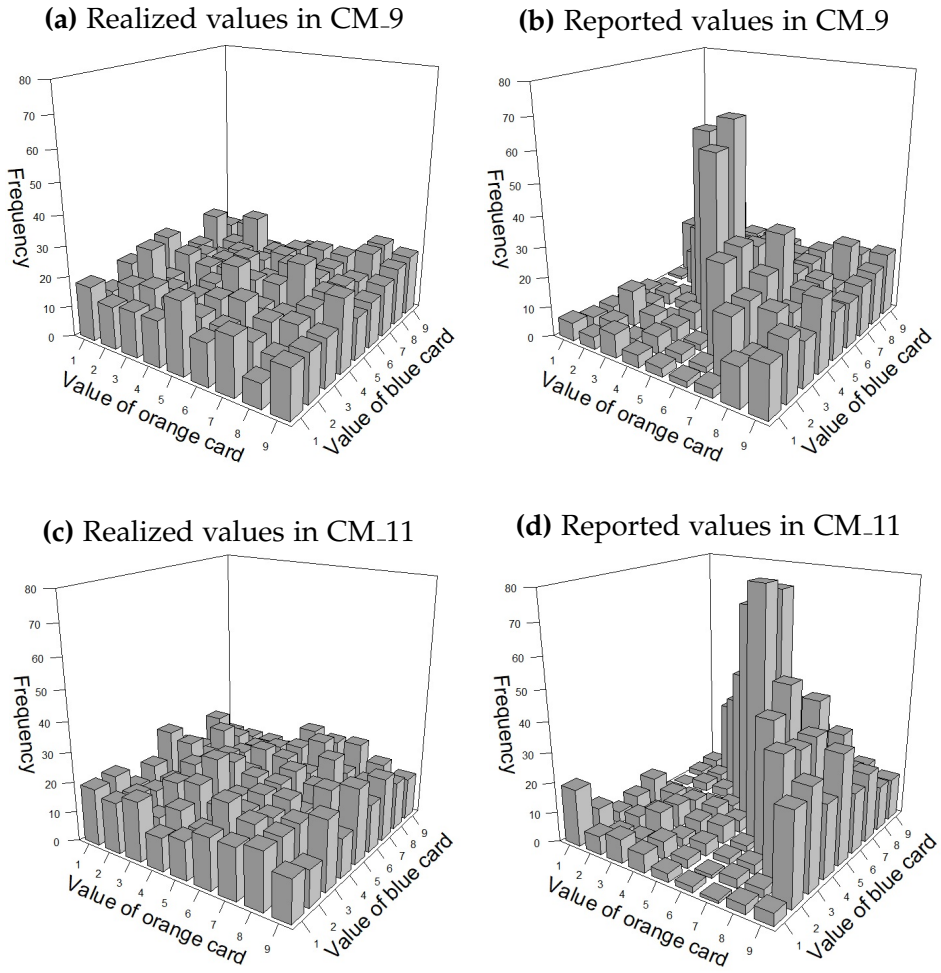


Figure 3.2: Frequency with which each combination of values was realized (left) vs. reported (right)

As expected, the realized values follow an approximately uniform distribution. Comparing this with the reported values, we observe a significant difference between the two distributions ($p - value < 0.001$ for both T levels, Kolmogorov-Smirnov test). Specifically, we note a concentration of the reported values within the area where the cards add up to at least the threshold of 9 or 11.

Furthermore, we notice an interesting pattern in senders' reporting strategy: values that add up to 9 or 11 are significantly more often reported than other higher values. The GR equilibrium does not make a sharp prediction in what the reporting strategy is concerned. However, assuming that receivers do not ignore the message, there are certain message strategies expected to decrease

the likelihood that the receiver will accept. One such example is when senders' misreports are clustered around a value (in our case the value is T). This allows the receiver to detect more easily bad hands misreported as good ones.

Next, if the reported hand is bad, receivers are sure to reject. Therefore, if senders are trying to maximize the receivers' acceptance probability, we would expect them to always report good hands. Nonetheless, if the highest value of the hand is smaller than 6 in CM_9 and 7 in CM_11, senders can be indifferent between reporting the bad hand truthfully or misreporting it as a good hand. This is because theoretically, receivers would reject in any case if they observe a value smaller than 6 (7). However, for bad hands where the highest value is 6 or 7 in CM_9 and 7, 8 or 9 in CM_11, to ensure the receiver accepts, senders should always misreport the lower card such that the reported sum is at least as high as T.

Figure 3.3 presents the proportion of bad hands misreported as being good for each value of the highest observed card both for CM_9[6] and for CM_11.

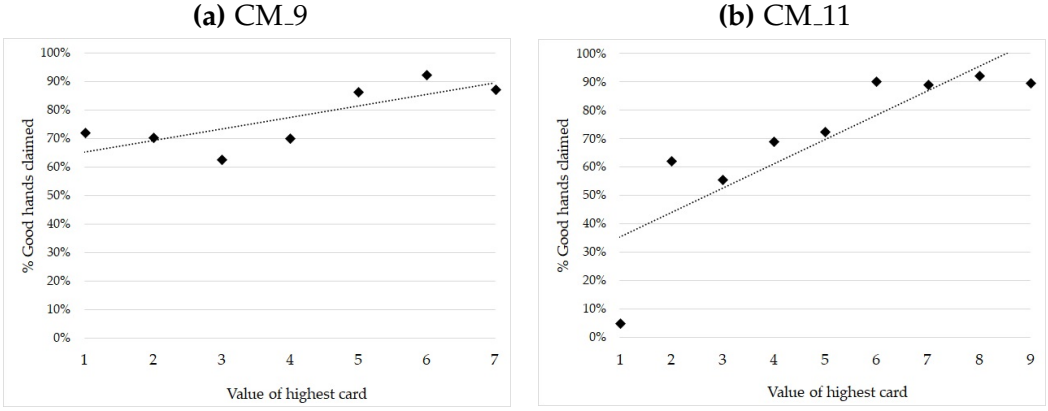


Figure 3.3: Proportion of bad hands reported as good

Pooling together the observations for values of 6 and 7 in CM_9 and 7, 8 and 9 in CM_11, we note that senders almost always (89% of the time in CM_9, and 90% in CM_11) report the observed bad hand as being good. We also note that senders are not indifferent between reporting truthfully or misreporting when the value of the highest card is lower than 6 (7). In fact, we observe a preference for misreporting bad hands as good when the highest card has a value of at most 5 in T=9. This holds for T=11 as well when the highest value is between 2 and 6, but not when this is 1.

Another interesting observation is the increasing trend - admittedly sharper in CM_11 than in CM_9 - in the proportions depicted in figure 3.3. This supports the idea that the stronger the evidence, the more compelled senders are

6 For CM_9, the maximum value of the highest card is 7 for this analysis as there are no bad hands that contain 8s or 9s.

to misreport bad hands as good. A rational explanation for this could be that the probability that a claim of having a good hand is believed increases with the value of the observed evidence[7].

From this analysis the following result ensues:

Result 1. *Senders in the CM_ treatments most commonly report bad hands as being good. This strategy is especially prevalent for hands containing at least a 6 in T=9 and at least a 7 in T=11.*

3.4.2 *Revelation strategy*

Figure 3.4 presents how often senders reveal the higher card[8] in the low conflict treatments (panel (a)) and in the high conflict ones (panel (b)).

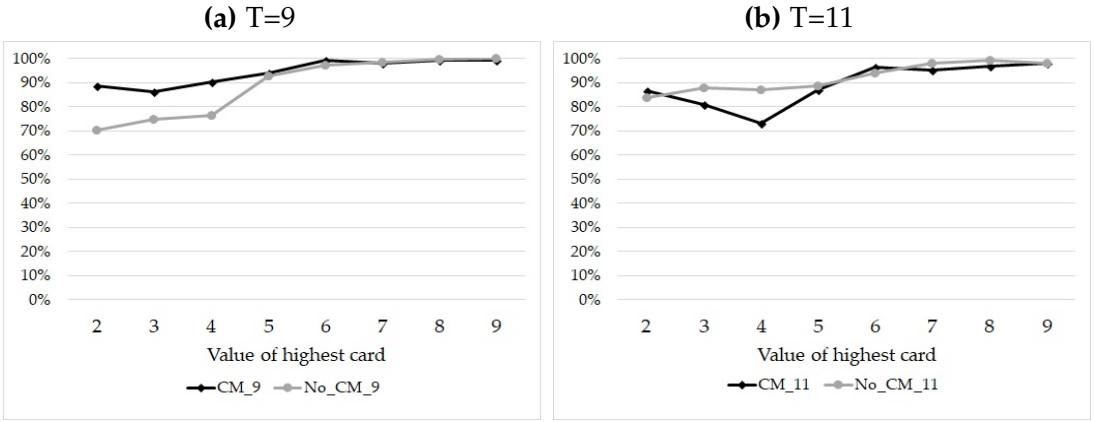


Figure 3.4: Percentage of times senders reveal higher card

From these plots we note that for values above 6 when T=9 and above 7 when T=11 senders reveal the higher card almost always (98.9% of the time in CM_9, 99.13% in No_CM_9, 96.93% in CM_11 and 98.53% in No_CM_11). There is no significant difference in terms of the revelation strategy for these values between the CM_ and No_CM_ treatments ($p - value = 0.407$ for T=9 and $p - value = 0.422$ for T=11, one-sided Wilcoxon signed rank test).

When looking at the cases where the value of the highest card is strictly less than 6, we note that revealing the higher card is still the most commonly used strategy (91% of the time in CM_9, 83% in No_CM_9, 86.88% in CM_11, 89.59% in No_CM_11). The frequency in the CM_9 treatment is weakly significantly higher than that in the No_CM_9 treatment ($p - value = 0.078$), while for

7 This assumes that the higher valued card is also the one that senders reveal, which is indeed the case as shown in Section 4.2 .

8 This analysis excludes the cases where the two cards had equal values.

the T=11 treatments, this difference is not significant ($p - value = 0.688$). We formulate the following result:

Result 2. *The higher card is revealed almost always in all treatments and this is not affected by messages.*

How does the strategy used by senders compare with equilibrium predictions? First, we need to note that GR's predictions are strict exclusively for the case where only one card is at least as high as 6 (7), while the other is strictly smaller than 6 (7). In this case, senders should always reveal the higher card. For the other two cases, when both are or none is higher than 6 (7) senders are indifferent between revealing the higher or the lower card.

Focusing on the case in which *only one card is at least as high as 6 (7)*, we find that senders reveal the higher card in 99.19% of cases in CM_9, 98.71% in No_CM_9, 97.06% in CM_11 and 98.92% in No_CM_11. Comparing the CM_9 and No_CM_9 treatments, we find no significant difference between the specified frequencies ($p - value = 1.000$), nor between each of these and the GR predictions ($p - value > 0.100$ for both CM_9 and No_CM_9).

There is no significant difference between the corresponding frequencies in the CM_11 and the No_CM_11 treatments either ($p - value = 0.313$). However, in these treatments, the proportion with which senders reveal the higher card when this is at least 7 is weakly significantly different from GR predictions ($p - value = 0.032$, one-sided Wilcoxon signed rank test). Nonetheless, the GR equilibrium seems to be a good predictor of the senders' revelation behavior.

3.4.3 *Acceptance strategy*

Figure 3.5 depicts the frequency with which receivers chose to accept for a given value of the card revealed by the sender (solid bars) across treatments. The graphs also plot how often receivers observed each value (hashed circles).

First, we notice that when messages are absent (right panels), receivers behave in line with equilibrium prediction: almost always rejecting hands when they observe a card below 6 (7), and accepting those when they observe a value above 6 (7). We use the same hit rate measure as in chapter 2 to analyze how often an observed outcome was predicted by the GR strategy. The hit rate in the No_CM_9 treatment is equal to 94.32% , while in No_CM_11 this is equal to 90.42%. When messages are present, we notice a significant decrease in hit rates for both treatments ($p - value = 0.008$). Specifically, in CM_9, the hit rate is equal to 85.76%, while in CM_11 it is equal to 82.64%.

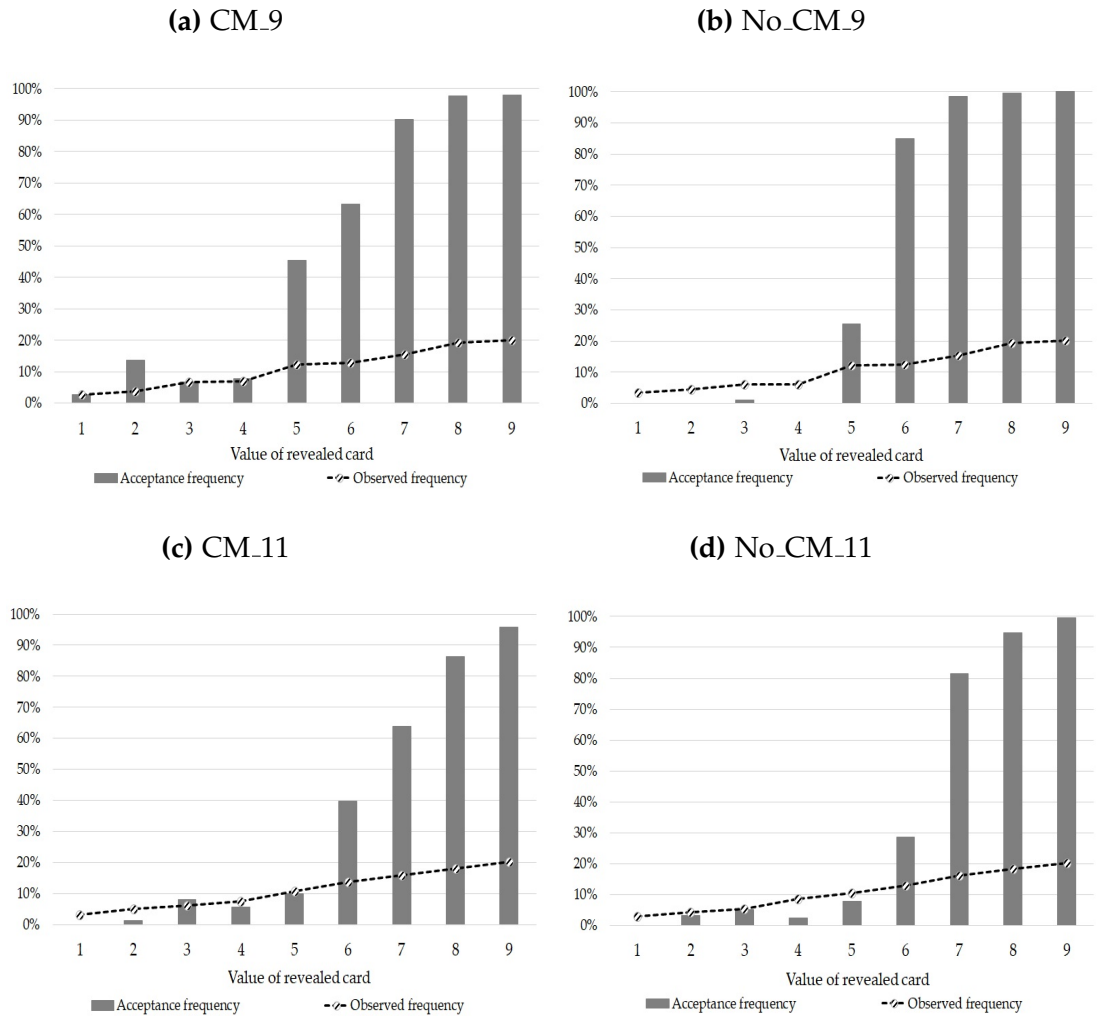


Figure 3.5: Acceptance and observed frequencies conditional on the value of the revealed card

Another way to visualize the deviations from the GR strategy is presented in figure 3.6. This depicts the distribution of errors across the entire state space for all treatments (CM_ in panels (a) and (d); No_CM_ in panels (c) and (f)) and for the GR strategy (T=9 in panel (b) and T=11 in panel (e)). We call an error the case of accepting a bad hand or rejecting a good one. Since the GR prediction accepts all hands that contain a card at least as high as 6 (7) and rejects everything else, the errors that this strategy gives rise to are clustered in well-defined areas. These are represented by bad hands containing a card with a value of at least 6 (7) and good hands where both cards are less than 6 (7). To construct the GR error distribution, we simulated the GR strategy on the realized state space. From figure 3.6 we see that the distribution of errors in the CM_ treatments is much more diffused across the entire state space. In the No_CM_ case we notice that errors are more concentrated around the same states as in the GR cases.

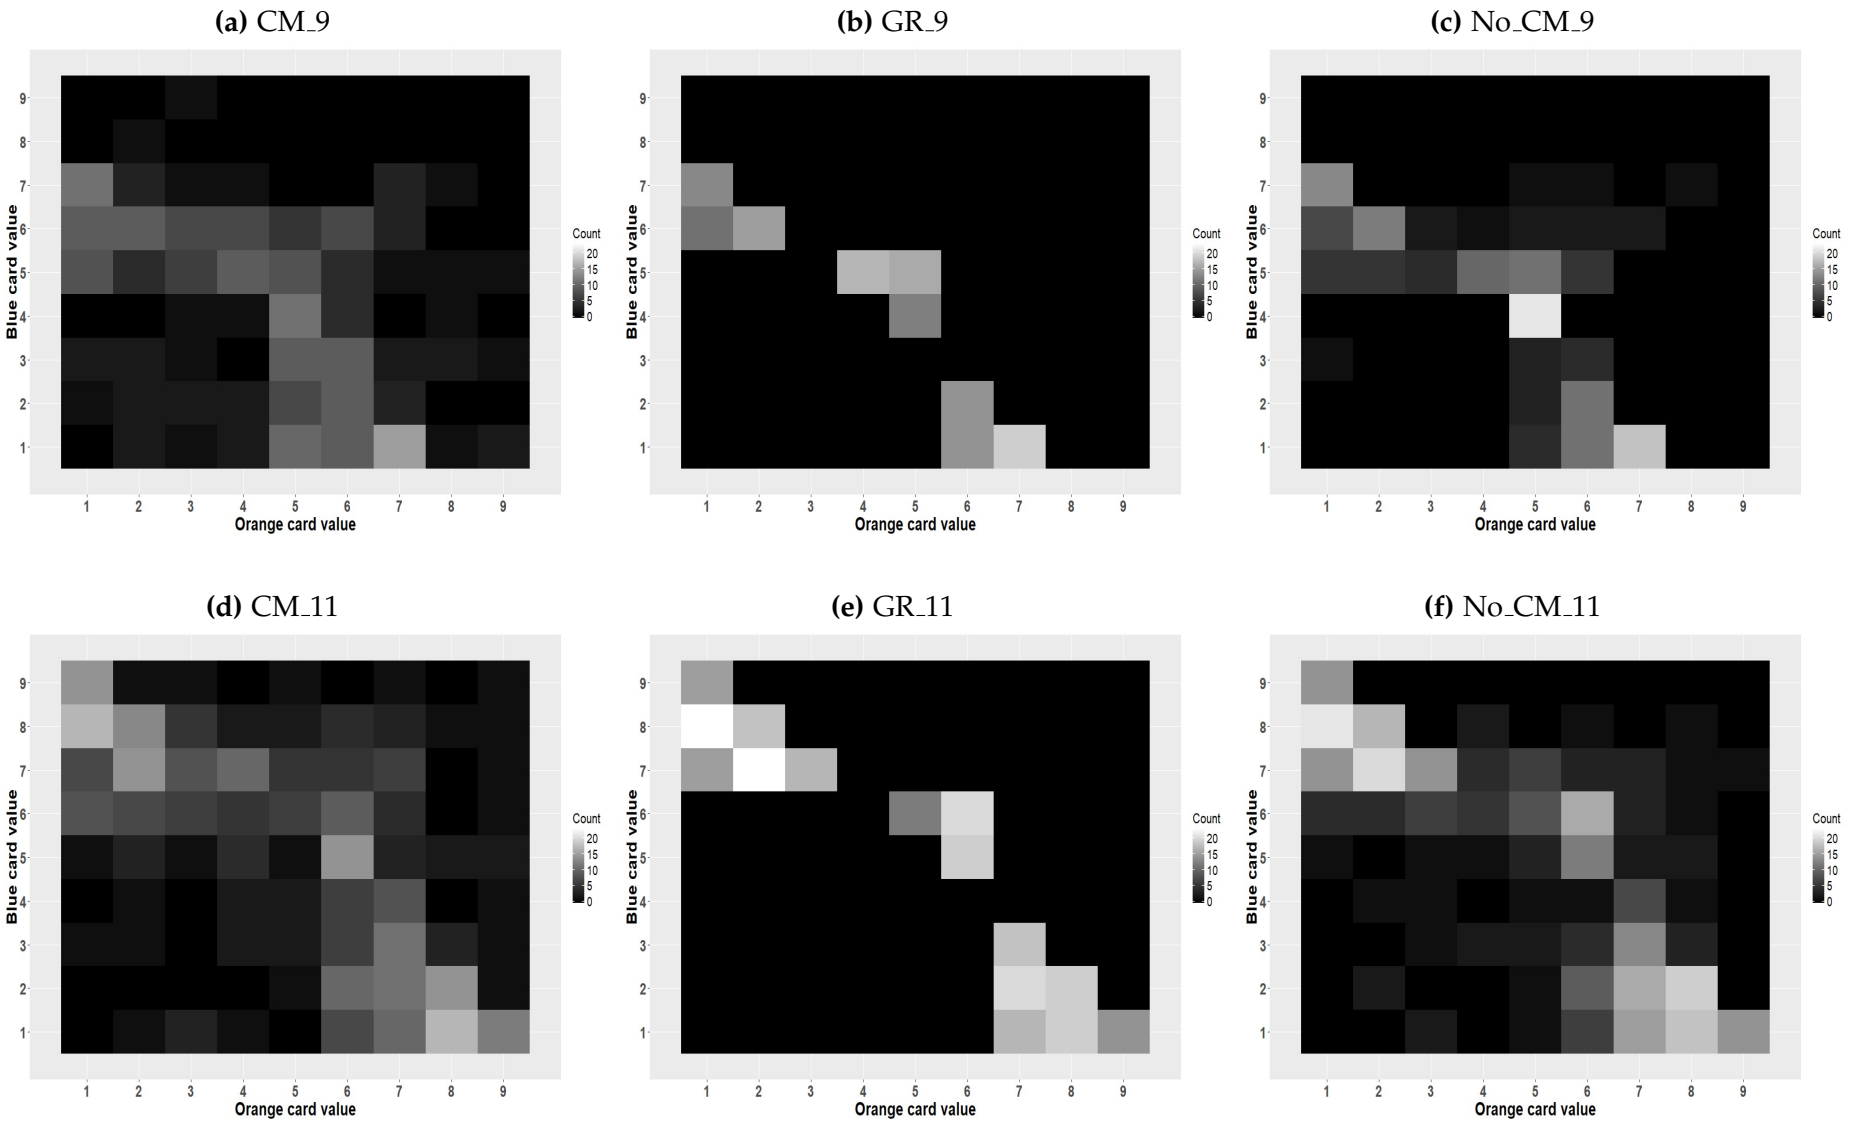


Figure 3.6: Distribution of errors across the state space (x-axis: values of the orange card; y-axis: values of the blue card)

This is clear also when looking at the left panels in figure 3.5 where the acceptance frequency plots show that receivers in the CM_ treatments tend to accept more often when the observed values are below the 6 (7) threshold and reject more often when they are above. These differences are significant for both conflict levels (see table 3.3) which points to the robustness of this effect.

Table 3.3: Acceptance frequencies across treatments

	T=9			T=11		
Observed value	CM_	No_CM_	p-value	CM_	No_CM_	p-value
$< 6(7)$	22.01%	9.77%	0.016	16.19%	11.56%	0.078
$\geq 6(7)$	89.59%	96.74%	0.016	83.31%	92.67%	0.023
Overall	67.73%	68.41%	0.742	52.50%	56.11%	0.102

This finding cannot be attributed to different distributions of observed values between treatments. Since senders almost always reveal their highest card in all treatments, the distribution of observed values (depicted with hashed dots in figure 3.5) is very similar across the CM_ and No_CM_ treatments. Result 3 follows:

Result 3. *Receivers accept values below 6 (7) and reject values above 6 (7) more often when communication is present.*

Pooling all hands, we do not observe a significant effect of messages on overall acceptance frequencies. This is because the deviations from the GR strategy for low and high values cancel each other out. Nonetheless, Result 3 is evidence against Hypothesis 1 as behavior is affected by messages (even though sender's behavior is not influenced by them - Result 2).

Lastly, we note that there is no learning in our experiment. Receivers' acceptance strategy in the first 15 rounds of the experiment does not differ from the one in the last 15 rounds. Neither does sender's revelation strategy in the first 15 rounds. Consequently hit rates are constant throughout the first and second half of the experiment.

3.4.4 *Average payoffs*

Table 3.4 presents the average payoffs of players across treatments.

Result 4. *Players' payoffs are consistently lower when messages are present; significantly so for the receiver in the low conflict case.*

Table 3.4: Average payoffs: CM_ vs. No_CM

	T=9			T=11		
	CM_	No_CM_	p-value	CM_	No_CM_	p-value
Sender	0.678	0.684	0.742	0.525	0.561	0.102
Receiver	0.830	0.879	<0.010	0.767	0.770	0.640

We see that contrary to Hypothesis 2, the receiver's payoff is negatively affected by communication. This is surprising given how messages can ultimately be ignored by the receiver. Nonetheless, it is in accord with Result 3 which suggests that messages are not ignored.

Why is the payoff effect not significant in the high conflict case? This seems surprising given that there are more deviations from GR predictions in CM_11 than in CM_9. However, receiver's acceptance behavior in No_CM_11 is also further from GR predictions than in No_CM_9 (recall that the hit rate was 90.42% in the former treatment and 94.32% in the latter one). This reduces the hit rate gap between the No_CM_ and CM_ treatments and consequently the effect of messages on receiver's average payoff. One explanation is that the T=11 game is more complex than the T=9 one since only one value (1) is certain to represent a bad hand. In the T=9 case, two values (8 and 9) were sure to be observed in case of a good hand. The higher degree of complexity can lead to more noisy behavior.

3.4.5 *Discussion*

From the previous section we note that both our hypotheses were rejected: messages affect behavior when partially informative evidence is present and this in turn affects final outcomes (equivalent to average payoffs). Two questions are consequently noteworthy:

1. Why are messages not ignored by receivers?

2. Why do receivers accept low values more often when communication is present than when it is absent, while also rejecting high values more frequently?

Regarding the first question, one explanation is that individuals simply cannot abstract from information they are faced with, be it relevant or not. This may be due to an "overimitation" bias which seems to be specific only to human social learning. This has been reported in a cognitive psychology study of Johnston et al. (2017) who find that certain canine species (shown to ex-

hibit human-like social learning characteristics) are better than children at ignoring irrelevant pieces of information when faced with solving a puzzle. It could also be due to a potential experimenter demand effect according to which receivers in our experiment think that they are supposed to take messages into account since they are always presented to them and senders are always obliged to send one.

A second possibility is that in our experiment messages are indeed useful and receivers take advantage of them. From the analysis of the reporting strategies we notice that senders do not randomize perfectly when misreporting. In fact, the reported values add up to T or T+1 too often: 41% of all reported good hands in CM_9 and 55% in CM_11. Receivers could then exploit this pattern by rejecting more often whenever the reported sum is equal to the threshold value. The probit regression in chapter 2 shows that receivers are indeed less likely to accept a reported good hand when the claimed sum is equal to T. Nonetheless, this departure of senders from babbling behavior does not justify the weight receivers put on messages and they would be better off behaving as the GR equilibrium predicts.

With respect to the second question, a tentative explanation for the skepticism part of receivers' behavior is offered by Ert et al (2014). The authors claim that the mere presence of communication emphasizes senders' persuasion incentives which makes receivers more reluctant to pursuing the action that is in senders' interests. This, however, would not explain why receivers tend to act against senders' interests when they become aware of their motives. Moreover, it does not clarify why the effect of communication is asymmetric depending on whether the observed value is high or low.

Another explanation could be that messages make receivers perceive the action of accepting as less risky than it actually is when the observed card value is small ($< 6(7)$), in which case it is more likely that a message reporting a good hand is false. On the other hand, when the observed value is high ($\geq 6(7)$) and it is more likely that the message is true, receivers might overweight the likelihood that the message is false. This could be a bias related to risk perception and further research is welcomed in this regard.

3.5 FOLLOW-UP TREATMENTS

To address the concern related to a potential experimenter demand effect we ran two additional follow-up treatments: R_choice and S_choice. These were variants of CM_9 were we observed a significant payoff difference.

In the R_choice treatment, we gave receivers the possibility to actively ignore the message. We implement this by giving them the option to not observe the message through choosing between "Display message" or "Don't display message" after the sender had submitted his reported values. In the S_choice treatment, we gave senders the possibility to not send a message by asking them to choose between "Send message" and "Don't send message". This part was introduced before the stage where they would have to report the values of the observed cards.

We implemented 3 sessions for these treatments, where the distribution of the realized card values corresponded to the distribution in the first 3 sessions of the CM_9 treatment. There were 144 participants equally distributed across treatments and sessions (24 participants in each session). 65.3% of hands were good (similar to the percentage when considering all 4 sessions - 65.8%). The same protocol as in CM_9 was followed. For the following analysis we will therefore consider only the data from the first 3 sessions in CM_9 and No_CM_9. Note that all previous results hold when taking into account the first 3 sessions only.

3.5.1 *Results*

3.5.1.1 *R_choice*

We begin the analysis by checking whether the significant negative effect of communication holds when receivers have more control over their information space. Table 3.5 presents the average payoff comparison across the relevant treatments.

Table 3.5: Average payoffs: R_choice, No_CM_9, CM_9

		R_choice vs No_CM_9		R_choice vs CM_9	
	R_choice	No_CM_9	p-value	CM_9	p-value
Sender	0.678	0.678	1.000	0.673	0.844
Receiver	0.863	0.882	0.406	0.831	0.156

We notice that the average payoffs of players in R_choice are higher than in CM_9, though not significantly so. Nonetheless, they are high enough to render the negative effect of communication insignificant (see comparison with No_CM_9). Note that the receiver's average payoff is significantly lower in CM_9 compared to No_CM_9 also when considering only the first 3 sessions ($p - value = 0.031$).

Result 5. *Messages have no effect when receivers have the option to not display them.*

In light of Result 4, Result 5 suggests that giving receivers more control over their information set has a positive effect on their average payoff. This is though not because they chose to use this option. In 98% of cases receivers choose to observe the message; hence, the information structure between CM_9 and R_choice is almost identical. Moreover, the result holds also when excluding the cases where there receiver did not look at the sender's message.

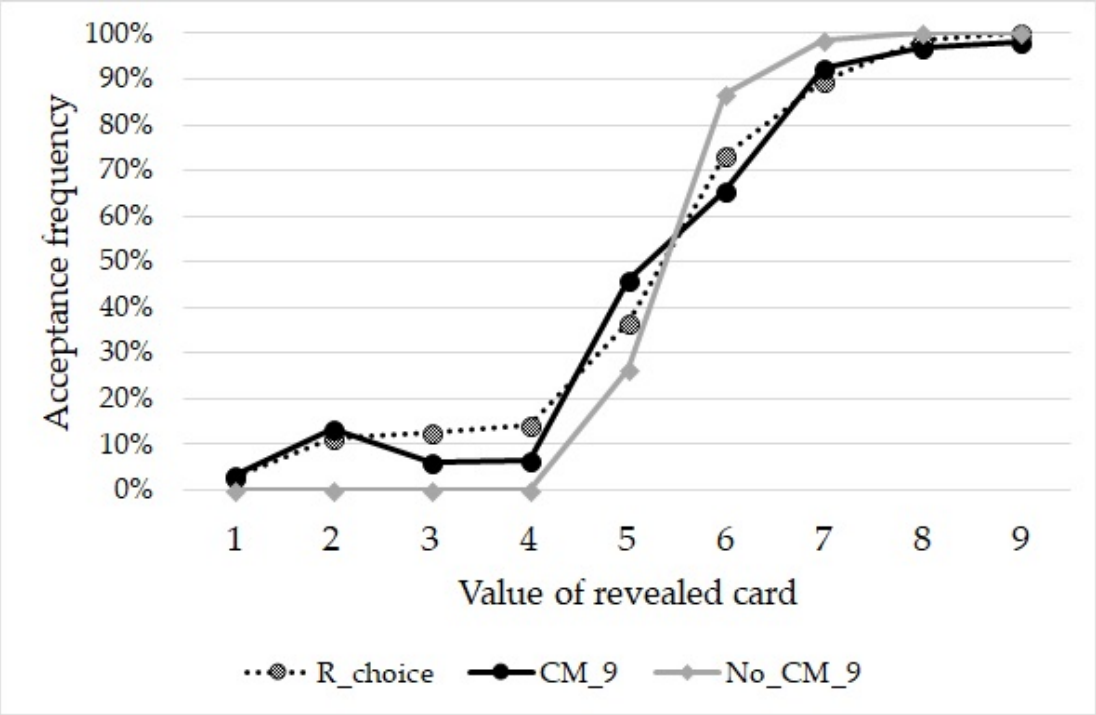


Figure 3.7: Acceptance frequencies: R_choice vs. CM_9 vs. No_CM_9

A potential explanation for the reduction in the negative effect of communication is the fact that senders misreport bad hands less often in R_choice compared to CM_9 (69.19% vs 76.55%). This difference is though not significant. The overall acceptance frequencies in R_choice and CM_9 do not differ either (see figure 3.7).

However, the hit rate in R_choice (87.35%) is higher (though not significantly) than in CM_9 (86.18%). Despite this, it still trails the 94.32% one in No_CM_9 ($p - value = 0.031$). This is due to receivers in R_choice accepting less often than those in CM_9 when the observed value is equal to 5 and more often when this is equal to 6. These differences lead to more outcomes being equivalent to the GR predictions in R_choice compared to CM_9. Moreover, the differences in the acceptance strategies for these values have a larger impact on hit rates and final payoffs as these values are observed more often than values equal to 3 or 4, where behavior in R_choice is further from GR predictions. Overall, average payoffs are relatively higher in R_choice than in CM_9. This is because senders are more truthful and receivers' acceptance strategy is closer to the GR predictions.

3.5.1.2 *S_choice*

Unlike in R_choice, where receivers have more control over their information space, in S_choice, receivers' control is reduced; in fact, more so than in CM_9. Recall that Result 5 suggested that the more control receivers have, the more likely it is that the negative effect of communication is reduced. Thus, in light of Result 5 one would expect that the payoff gap due to messages (Result 4) should be widened in the S_choice case. However, if senders prefer to stay silent rather than lie, one could also observe a more informative message strategy, allowing receivers to better separate between good and bad hands.

Table 3.6 presents the payoff comparison between the S_choice follow-up treatment and the two main treatments (data from first 3 sessions only). We notice that indeed, the direction of the effect of communication is the predicted one. As opposed to the R_choice case, average payoffs are lower in S_choice compared to No_CM_9 (the difference is significant at 10% level). Compared to CM_9 though, average payoffs are lower but not significantly so.

Table 3.6: Average payoffs: S_choice, No_CM_9, CM_9

		S_choice vs No_CM_9		S_choice vs CM_9	
	S_choice	No_CM_9	p-value	CM_9	p-value
Sender	0.645	0.678	0.438	0.673	0.844
Receiver	0.829	0.882	0.063	0.831	0.438

Result 6. *Messages have a (weakly) significant negative effect on receivers' average payoff when senders have the option to not send a message.*

Result 6 is in line with our finding that decreasing receiver's control over his information space harms the receiver. To see why the option to not send a message does not lead to more informative messages we note that in the 18% of cases where no message was sent, 49% of these were good hands. Therefore, sellers did not use this option to refrain from lying. Because of this, Result 6 holds also when considering only the cases in which messages were sent in S_choice. Moreover, the percentage of misreports is not significantly different in S_choice compared to CM_9. It is not surprising therefore that the receivers' acceptance strategy in S_choice is closer to that in CM_9 rather than the one in No_CM_9, as can be observed in figure 3.8. This is confirmed also by comparing the hit rates across treatments: the rate in S_choice is 86.18%, which is identical to the one in CM_9.

Perhaps the most striking difference between the acceptance frequencies in S_choice compared to those in No_CM_9 is the pronounced skepticism for values at least as high as 6. These are accepted in 89% of cases in S_choice which

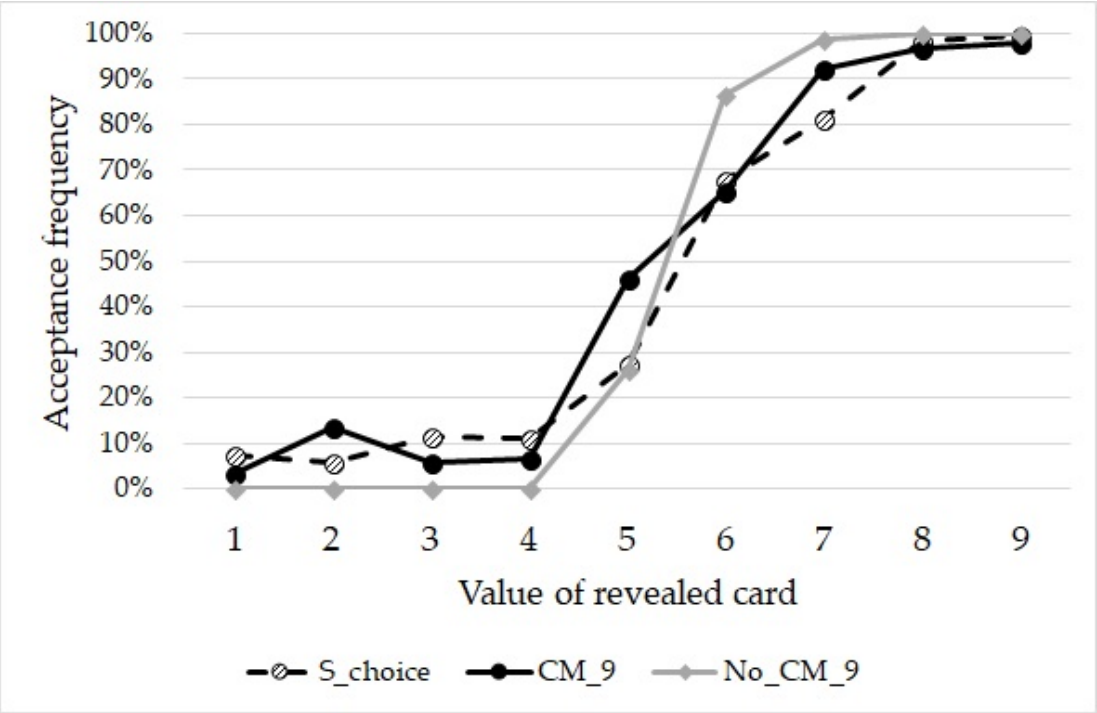


Figure 3.8: Acceptance frequencies: S_choice vs. CM_9 vs. No_CM_9

is significantly lower than 97.19%, the frequency in No_CM_9 ($p - value = 0.031$). The acceptance frequencies in S_choice and No_CM_9 for values less than 6 are not significantly different.

3.5.2 *Discussion*

In the follow-up treatments we control for the possibility that certain features of the experimental design might make receivers believe that they are expected to take messages into account. In this process we implemented two treatments that vary in the degree of control the receiver has over his information set.

We find an interesting dynamic arising by simply giving the option to the receiver to actively not look at the message. Even though they almost always choose to see the message, receivers tend to make fewer errors in R_choice than in CM_9 where the choice was not available. We propose two explanations for this effect.

On the one hand, senders become more truthful. This can very well be strategic as truth-telling can signal to the receivers that it is worth looking at the message, which is what senders believe might increase the acceptance probability. Another possibility is that telling a lie no longer pays off if the receiver might not look at it. Irrespective of what makes senders more truthful,

it is a fact that this helps receivers to better discern between good and bad hands.

On the other hand, receivers are less skeptical about high values in R_choice compared to CM_9. This can be a result of the trustworthiness that senders' more truthful behavior might instill. It could also be that more control over one's information space strengthens the belief that deception can be more easily detected. This helps receivers make better assessments of the likelihood that the hand is good given the observed card value.

Giving the choice to senders regarding whether to send a message or not reduces even further the control that receivers have over their information space compared to CM_9. From the perspective of the previous argument about the effect of control over one's information space, we can understand why this treatment does not improve receivers' decision-making strategies. It is worth noting that neither player is worse off either.

3.6 CONCLUSION

We investigated the effect of communication on the interpretation of partially-informative evidence. Using a sender-receiver game with asymmetric information where both cheap talk and hard evidence were present, we found that cheap talk can have welfare reducing effects by significantly lowering receiver's average payoff while not affecting the sender's one. This is because receivers are unable to ignore senders' messages which distort their interpretation of evidence. In a follow-up investigation where we increase receivers' control over their information space we reconcile the finding from chapter 2 that such a manipulation can improve receivers' inferences from the objective information received.

To the best of our knowledge, this is the first study in this literature pointing to a clearly negative effect of communication. To examine the pervasiveness of this phenomenon, we would welcome future research trying to disentangle its causes. At this point, our study can be used to raise awareness among consumers regarding the types of buyer-seller interactions they might engage in and how they could alter the design to improve the outcome of the transaction. Political campaigns are one other environment where cheap talk and evidence abound. It is unclear though to what extent voters control the structure of their information set. Understanding this can lead to better informed voting decisions.

BIBLIOGRAPHY

Cai, H. and Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36.

Charness, G. (2000). Self-serving cheap talk: A test of Aumann's conjecture. *Games and Economic Behavior*, 33(2):177–194.

Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6):1579–1601.

Charness, G. and Dufwenberg, M. (2011). Participation. *American Economic Review*, 101(4):1211–37.

Clots-Figueras, I., Hernán-Gonzaléz, R., and Kujal, P. (2015). Information asymmetry and deception. *Frontiers in Behavioral Neuroscience*, 9(109).

Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Ert, E., Creary, S., and Bazerman, M. (2014). Cynicism in negotiation: When communication increases buyers' skepticism. *Judgment and Decision Making*, 9:191–198.

Forsythe, R., Lundholm, R., and Rietz, T. (1999). Cheap talk, fraud, and adverse selection in financial markets: Some experimental evidence. *The Review of Financial Studies*, 12(3):481–518.

Glazer, J. and Rubinstein, A. (2006). A study in the pragmatics of persuasion. *Theoretical Economics*, 1:395–410.

Hagenbach, J. and Perez-Richet, E. (2018). Communication with evidence in the lab. *CEPR Discussion Paper No. DP12927*.

Jin, G. Z., Luca, M., and Martin, D. (2015). Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure. *NBER Working Paper No. 21099*.

Johnston, A. M., Holden, P. C., and Santos, L. R. (2017). Exploring the evolutionary origins of overimitation: a comparison across domesticated and non-domesticated canids. *Developmental Science*, 20(4).

Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821.

Mohlin, E. and Johannesson, M. (2008). Communication: Content or relationship? *Journal of Economic Behavior & Organization*, 65(3-4):409–419.

Penczynski, S. P. and Zhang, S. (2016). Disclosure of verifiable information under competition: An experimental study. *Mimeo*.

Roth, A. E. (1995). *Bargaining experiments*. Princeton University Press.

Rubinstein, A. (2012). *Economic Fables*. Open Book Publishers, Cambridge, UK.

Zultan, R. (2012). Strategic and social pre-play communication in the ultimatum game. *Journal of Economic Psychology*, 33:425–434.

4

COMMUNICATION, SKEPTICISM AND LYING AVERSION

4.1 INTRODUCTION

This chapter explores the role of communication in sender-receiver games without evidence but with information asymmetry and sender state independent preferences. These games are of particular appeal as they model many economic situations such as investment, voting or purchase decisions. Consider the example of a seller who would like the buyer to buy his product irrespective of its quality or a politician who always wants the voter to prefer him rather than the opponent candidates. Such preferences have been given limited attention in the cheap talk literature due to the result that informative equilibria cannot arise in this setting (Crawford and Sobel, 1982). In light of experimental findings suggesting that communication has strong behavioral effects, we are interested in identifying the extent to which communication can affect outcomes in these settings.

The game we focus on involves a seller endowed with two products, red and blue, and a buyer. The state of the world is binary, red or blue, and the distribution of these states is common knowledge. If the state is blue, the blue product has a higher quality than the red one; if the state is red, the red product is of higher quality. The seller observes the randomly determined state and sends a message to the buyer regarding its color. The buyer observes the message and decides whether to buy the red product or the blue one. Preferences are such that the seller would like the buyer to always buy the red product. However, the buyer would like to buy the product of higher quality. Therefore, the buyer's preferences depend on the state, while the seller's do not. According to standard theory, communication should have no effect on final outcomes in this setting.

Ert et al. (2014) implement a conceptually similar environment and find that the mere existence of cheap talk makes the receiver skeptical about choosing his optimal action. We observe a similar effect in chapter 3 though the settings are not directly comparable. Ert et al. (2014) explain this result by stating that communication makes receivers aware of the persuasive intentions of the senders which leads receivers to become more cynical and act conservatively. This is in line with Darke and Ritchie (2007) and Friestad and Wright (1994) who suggest that consumers could react more cautiously when they are in-

formed about advertising tactics. These findings suggest that communication can lead to worse outcomes for both parties as it can induce skepticism to receivers and lead them to reject good options they would otherwise accept.

On the other hand, a wide body of literature regards communication as a tool to improve economic outcomes. Specifically, it has been found to increase donations in dictator games and agreement frequencies in ultimatum games (Mohlin and Johannesson, 2008; Roth, 1995), to increase trust in trust games (Charness and Dufwenberg, 2006, 2011) and coordination rates in coordination games (Cooper et al., 1992). It has also been established that it allows players to outperform the expectations of economic models in auctions[1] (Bazerman et al., 1998; Valley et al., 2002). Furthermore, in social dilemmas such as modified dictator or public goods games, communication consistently improves cooperation rates (Bochet et al. (2006), Bicchieri and Lev-On (2007) for an overview). The size of these effects can depend on the communication protocol: one-way or two-way (Burton et al., 2005; Charness, 2000), face-to-face or computerized (Bicchieri and Lev-On, 2007). In sender-receiver games, lying aversion (Abeler et al., 2016; Kartik, 2009) could lead to positive effects from communication. In fact, over-communication (transmitting more information than theory would predict) has been empirically documented in such games (Cai and Wang, 2006; Sánchez-Pagés and Vorsatz, 2007).

Using a simple theoretical analysis we show that depending on the distribution of the states of the world, outcomes are driven by either skepticism or lying aversion. Therefore, communication can have both positive and negative effects in our game. Specifically, when the red state is more likely than the blue one, communication can harm due to receiver's increased skepticism. We hypothesize that the mechanism behind this effect is an aversion to being profitably deceived, which we call *"sucker aversion"*. On the other hand, when the blue state is more likely than the red one, the lying aversion effect takes over, reducing buyer's skepticism and making both parties better off. Note that the distribution of states represents the likelihood that the seller's preferred action coincides with the buyer's optimal one.

We implement a lab experiment to test these hypotheses using the one-shot seller-buyer game described above. We use four treatments across which two dimensions are varied: whether the seller is allowed to send a free-form text message to the buyer and the frequency of red balls in the urn. The red ball frequency can take two values: $\frac{1}{3}$ and $\frac{2}{3}$.

Our results show that communication affects outcomes. Specifically, our findings support the lying aversion hypothesis as communication significantly reduces skepticism in the $\frac{1}{3}$ - treatments while significantly increasing seller's

1 This refers to the prediction that the Chatterjee-Samuelson equilibrium is the most efficient one of a double auction with two-sided private information.

average payoff. Nonetheless, this effect is inconsistent with the requirement of our model that $\frac{1}{2}$ of the sellers are lying averse since the observed fraction in our data is much lower. To this extent, buyers' behavior in the corresponding treatment without communication is revelatory. There, we find that buyers' behavior is significantly different than equilibrium predictions as buyers buy the red product more often than theory predicts. This is rather surprising given the simple environment we implement. Nevertheless, efficiency concerns could give rise to this behavior since buying the red product ensures a high payoff for the seller. Such preferences might, in turn, magnify the effect of lying aversion when communication is present. Furthermore, in contrast with (Ert et al., 2014), communication does not increase skepticism in the $\frac{2}{3}$ - treatments, and does not have a significant effect on players' payoffs. Therefore, we find no evidence for sucker aversion.

The rest of the chapter is structured as follows: Section 4.2 gives a more detailed overview of the relevant literature, Section 4.3 presents the standard theoretical analysis and the effect of introducing non-standard preferences, Section 4.4 presents the experimental design and procedures; this is followed by the results analysis in Section 4.5 while Section 4.6 discusses the results and concludes.

4.2 RELATED LITERATURE

This section focuses on experimental studies related to communication about private information. An important study in this field is that of Cai and Wang (2006). This paper implements treatments to test the predictions of the Crawford and Sobel (1982) model of strategic information transmission (cheap talk). In their game, players are paired and assigned to the role of sender or that of receiver. The sender has private information about the state of the world, s. The value of s is a number uniformly distributed over the state space $S = \{1,3,5,7,9\}$. They use strict-form messages which can take any value from $M = \{1,3,5,7,9\}$. After observing the message, the receiver takes an action from $A = \{1,2,3,4,5,6,7,8,9\}$. The receiver wants to take an action as close as possible to the state while the sender's preferred action is $s+d$, where d represents the preference difference (size of the conflict of interests). They implement treatments with different values for d and find support for the theoretical predictions according to which as d increases, less information is transmitted by senders and used by receivers. As opposed to equilibrium predictions, they observe that senders report the state truthfully more than they should (over-communication). This result is in line with Dickhaut et al. (1995) who investigate a similar topic in a similar environment. We differ from these papers in several ways. First, in our game the sender has state independent preferences and therefore the parameter d is not relevant. Second, we replace

the concept of deterministic preference difference with that of the likelihood that preferences differ by varying the receiver's optimal action (defined as such from a standard theoretical point of view). Third, we implement treatments also without communication so that we can compare behavior with an empirical benchmark as well, not only with a theoretical one.

Blume et al. (1998, 2001) also investigate the strategic transmission of information in an experimental setting and their results are largely in line with the over-communication phenomenon documented by Cai and Wang (2006), However, in their studies the focus is on how the meanings of messages evolve giving rise to a different setting than the one we are investigating in this chapter. Nevertheless, these results support the hypothesis that lying aversion is a robust phenomenon in strategic frameworks.

Another relevant paper is Forsythe et al. (1999) who investigate a buyer-seller setting characterized by inefficient trading due to adverse selection (Akerlof, 1970). They test whether two different communication mechanisms lead to efficiency gains (more trade). In the first mechanism, "cheap talk", sellers' message can be any subset of product qualities (not necessarily the true one). In the other mechanism, "antifraud", sellers' report must include the true quality. They find that both mechanisms increase trade efficiency but in the cheap talk environment this gain comes at the buyer's expense. Our studies are similar in that the seller's preferences are state independent, while the buyer's are not. Nevertheless, we differ in several dimensions. In their experiment, buyers and sellers with different valuations for an asset have to trade in a market by bidding prices. Therefore, unlike in our setting, sellers are active players even when communication is not possible. Also, they focus only on a framework where the different asset qualities are equally likely while we explore scenarios in which the distribution is not uniform.

The closest paper to our study comes from the negotiation literature where Ert et al. (2014) investigate behavior in a cheap-talk game with a seller who has state independent preferences. In their game, the seller randomly draws two cards from a deck containing 100 cards with consecutive values from 1 to 100. The buyer is then informed about the value of the lower card and decides whether to buy the two cards from the seller or not at a fixed price. The seller's payoff is the fixed price while the buyer's payoff is represented by the sum of the two cards minus the fixed price. In their experiments, the authors fix the value of the lower card to 40 which makes it optimal for the buyer to buy the cards[2]. They compare treatments with and without communication to investigate whether communication makes buyers more skeptical to buy and find a significant skepticism effect. In our study, we simplify the setting

2 They do so by using two rounds: one in which the lower card's value is 40, and one in which it is randomly drawn. The subjects are not informed about the value of the lower card but only that in one of the two rounds this value was pre-determined.

by implementing a cheap-talk game with only 2 possible states. This should reduce the noise that free-form communication might introduce. Second, we are interested in studying the evolution of skepticism when the buyer's optimal action is different than the seller's preferred action. In the terms of Ert et al. (2014) this is equivalent to the buyer's ex ante optimal action being not to buy.

Serra-Garcia et al. (2011) investigate the effect of vague and precise communication in a sequential 2-player public good game with asymmetric information about the value of the public good. This can be high, intermediate or low. Like us, they hypothesize that communication can decrease efficiency, though their mechanism is based on the leader's lying aversion rather than on the follower's skepticism. They do not find support for this hypothesis as leaders frequently lie with precise communication, while when messages can be vague, followers do not correctly interpret them.

Other similar game environments are found in the literature on deception detection. Belot and Van de Ven (2017) also use a buyer-seller setting with asymmetric information. The main difference with our study is the communication protocol. Specifically, Belot and Van de Ven (2017) are interested in whether buyers can spot deception in face-to-face interactions. We, on the other hand, are interested in whether anonymous communication can improve outcomes compared to a benchmark in which communication is not allowed. Nonetheless, the results of this study are relevant as the ability to spot deception can be a mechanism that leads to gains from communication. The authors find that buyers are better than chance in predicting which sellers are lying when they interact through face-to-face communication. This result supports the finding of Belot et al. (2012) who study the ability of people to detect deceit during a television game show where players are competing in a high-stakes, simultaneous prisoner's dilemma game. Other studies using face-to-face interaction also find that subjects are better than chance when asked to predict the behavior (Brosig, 2002; Dawes et al., 1977) or private information of others (Ockenfels and Selten, 2000). Deceit-detection in written communication, such is the case in our study, is though less explored in the literature. Chen and Houser (2017) investigate whether observers of a three-person trust game can predict the trustworthiness of free-form written messages. They find little support for subjects' ability to predict better than chance.

4.3 THEORETICAL ANALYSIS

In this section we first describe the game and the theoretical predictions according to standard theory. Then, we explore how these predictions might change if we consider lying averse sellers who incur a psychological cost from

lying and buyers who are averse to being profitably deceived. Such buyers incur a psychological cost from taking an action which rewards a deceiving seller. We will call this cost a sucker cost. Furthermore, we investigate the role of these behavioral biases when incentives are more likely to be aligned compared to when it is more probable that they are unaligned.

4.3.1 *The game*

In the benchmark case of our experiment, the following game is played between two players: a seller and a buyer. The timeline of the game is the following:

Stage 1: Nature's move. *Nature* randomly determines the state of the world (s), represented by a ball that can take one of two possible colors: red (R) or blue (B). Hence, the state space is $S = \{R, B\}$. The state s is private information of the seller. The probability of $s = R$ is p and this is common knowledge.

Stage 2: Seller's move. After observing the state, the seller can send a free-form message (m) to the buyer. For the purpose of this analysis we will restrict the message space to two messages: reporting that the ball is red and reporting that the ball is blue. The message space under consideration is therefore $M = \{R, B\}$.

Stage 3: Buyer's move. After receiving the message from the seller, the buyer decides whether to buy the red product or the blue product. This is equivalent to taking an action (a) from the following set: $A = \{R, B\}$.

The payoff of the seller depends only on the buyer's action while the buyer's payoff depends both on his action and on the true state of the world. Table 4.1 summarizes players' material payoffs.

Table 4.1: Payoff Matrix (seller's payoff listed first in each cell)

	$a = R$	$a = B$
$s = R$	(1, 1)	(0, 0)
$s = B$	(1, 0)	(0, 1)

The seller, therefore, would always like the buyer to buy the red product (state independent preferences) while the buyer would like to buy the red product only when the state is red and to buy the blue product when the state is blue. We assume that messages have literal meanings. In a series of experiments Blume et al. (2001) show that indeed, players tend to use messages with their natural language interpretations.

In our experiment, we implement different treatments where p takes one of the following two values: $p \in \{\frac{2}{3}, \frac{1}{3}\}$. Essentially, what p represents is the likelihood that the interests of the seller and the buyer are aligned. We chose these values to reflect a case where interests are more likely to be aligned ($p = \frac{2}{3}$) and one where they are more likely to be conflicting ($p = \frac{1}{3}$).

4.3.2 *Equilibrium analysis*

Our solution concept is Perfect Bayesian Equilibrium (PBE). We restrict our analysis to pure strategies. A strategy for the seller is a vector $m = (m(R), m(B))$ where $m(R)$ represents the seller's message decision when the state is red, and $m(B)$ his message decision when the state is blue. A strategy for the buyer is a vector $a = (a(R), a(B))$ where $a(R)$ represents the buyer's action when the message is red, and $a(B)$ his action when the message is blue. Moreover, let $\mu(m) \in [0, 1]$ be the posterior probability that the buyer assigns to $s = R$ upon observing m. All players are risk neutral, material payoff maximizers.

Note that in absence of messages the buyer is not indifferent between the two available actions. It is easy to see that the buyer's ex ante optimal action when $p = \frac{2}{3}$ is $a = R$, whereas when $p = \frac{1}{3}$ it is $a = B$. When introducing messages, it is important to distinguish between three types of equilibria: uninformative, informative and persuasive ones. In an uninformative equilibrium, the buyer's beliefs and actions are the same as when communication would not be possible. In an informative equilibrium the buyer's belief about the probability that $s = R$ conditional on the message is higher than the prior but not high enough to make the buyer choose an action different than his ex ante optimal one. Lastly, in a persuasive equilibrium, the buyer's optimal action is different than his ex ante one in a way that the seller's payoff is higher than in any uninformative equilibrium.

In a recent paper, (Lipnowski and Ravid, 2017) prove that in a cheap-talk game with two buyer actions, two states and sender state independent preferences where the buyer is not indifferent between the two actions at the prior, informative equilibria exist but persuasive ones are impossible. The following two lemmas follow trivially.

Lemma 1. *In every PBE of the game where $p = \frac{2}{3}$, seller's message strategy (m) is not persuasive and $a(R) = a(B) = R$, on the equilibrium path.*

Lemma 2. *In every PBE of the game where $p = \frac{1}{3}$, seller's message strategy (m) is not persuasive and $a(R) = a(B) = B$.*

Note that when $p = \frac{2}{3}$, it is relatively more likely that the interests of the players are aligned and the buyer's optimal action coincides with the seller's

preferred one. The opposite is true when $p = \frac{1}{3}$. Any seller strategy can be supported in equilibrium in both cases. However, if messages have a natural language interpretation (Blume et al., 2001), it is reasonable to assume that sellers would be more likely to signal their preferred action, i.e. red, when the state coincides with this action, rather than signaling that the state is blue.

4.3.3 *Introducing lying and sucker costs*

4.3.3.1 *Lying cost*

We begin with a simple analysis of players' behavior when sellers are potentially lying averse. We will call this game "the L-game". For simplicity, we assume that there are two types of sellers: liars (LI) and truth-tellers (TT). Liars are material payoff maximizers. Truth-tellers are assumed to be lying averse in that they incur a psychological cost from sending a message $m \neq s$. This cost is high enough such that they always tell the truth.

We implement this in our game by letting nature decide not only the state but also the seller's type. Let $l \in [0, 1)$ be the proportion of liars in the population of sellers. We assume that l is common knowledge. Therefore, the buyer does not know whether he is interacting with a TT or LI seller but he knows the distribution of these two types in the population of sellers. The case where $l = 1$ is equivalent to one in which sellers are selfish rational utility maximizers and the analysis is the same as in Section 4.3.2. When $l = 0$, all sellers are truth-tellers and the buyer's optimal strategy is to follow the message.

There are four possible (pure) strategies for the LI seller:

(a) $(m_{LI}(R), m_{LI}(B)) = (R, R)$

(b) $(m_{LI}(R), m_{LI}(B)) = (B, B)$

(c) $(m_{LI}(R), m_{LI}(B)) = (R, B)$

(d) $(m_{LI}(R), m_{LI}(B)) = (B, R)$

The TT seller always chooses $(m_{TT}(R), m_{TT}(B)) = (R, B)$. Next, we will check which of the potential LI seller strategies can be supported in a PBE of the two parameterizations of the game under consideration.

Case 1. $p = \frac{2}{3}$

(a)

Lemma 3. *The following strategies:* $(m_{LI}(R), m_{LI}(B)) = (R, R)$, $(a(R), a(B)) = (R, B)$, *are part of a PBE of the L-game where* $p = \frac{2}{3}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (R, R)$. Buyer's beliefs consistent with the specified seller strategy are:

$$\mu(R) = \frac{pr(m = R|s = R) \cdot pr(s = R)}{pr(m = R)} = \frac{1 \cdot \frac{2}{3}}{\frac{2}{3} + \frac{1}{3}l} = \frac{2}{2+l}$$

$$\mu(B) = \frac{pr(m = B|s = R) \cdot pr(s = R)}{pr(m = B)} = 0$$

Given these beliefs, the buyer best replies by maximizing his expected utility given each message.

$$EU_{buyer}(a = R, m = R) = 1 \cdot \mu(R) + 0 \cdot (1 - \mu(R)) = \frac{2}{2+l}$$

$$EU_{buyer}(a = B, m = R) = 0 \cdot \mu(R) + 1 \cdot (1 - \mu(R)) = \frac{l}{2+l}$$

$$EU_{buyer}(a = R, m = B) = 1 \cdot \mu(B) + 0 \cdot (1 - \mu(B)) = 0$$

$$EU_{buyer}(a = B, m = B) = 0 \cdot \mu(B) + 1 \cdot (1 - \mu(B)) = 1$$

This means that $EU_{buyer}(a = B, m = B) > EU_{buyer}(a = R, m = B)$. Moreover, since $l < 1$ we have that $EU_{buyer}(a = R, m = R) > EU_{buyer}(a = B, m = R)$. Therefore, the buyer's best reply strategy is $(a(R), a(B)) = (R, B)$.

Finally, does any LI seller type have an incentive to deviate given the buyer's strategy? Clearly, when $s = R$, no LI seller type would like to deviate by sending $m_{LI}(R) = B$ since that would lead to a payoff of 0, which is smaller than the payoff the seller gets from $m_{LI}(R) = R$. When $s = B$, $EU_{LI}(m = R, s = B) = 1$ while $EU_{LI}(m = B, s = B) = 0$, so the LI seller does not want to deviate from the original strategy $(m_{LI}(R), m_{LI}(B)) = (R, R)$.

□

For the following cases, buyer's beliefs and expected utilities from each action are computed like in Lemma 3 and the explicit formulas are omitted.

(b)

Lemma 4. *The following strategies:* $(m_{LI}(R), m_{LI}(B)) = (B, B)$, $(a(R), a(B)) = (R, R)$ *are part of a PBE of the L-game where* $p = \frac{2}{3}$ *if and only if* $l \geq \frac{1}{2}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (B, B)$. Buyer's beliefs consistent with this strategy are: $\mu(R) = 1$, $\mu(B) = \frac{2l}{2l+1}$. Consequently, his expected utilities from either choice are: $EU_{buyer}(a = R, m = R) = 1$, $EU_{buyer}(a = B, m = R) = 0$, $EU_{buyer}(a = R, m = B) = \frac{2l}{2l+1}$, $EU_{buyer}(a = B, m = B) = \frac{1}{2l+1}$.

Therefore, the buyer best replies by choosing $a(R) = R$ irrespective of l. When $m = B$, his optimal choice depends on l as follows. If $l \geq \frac{1}{2}$, the buyer best replies by choosing $a(B) = R$. This is an equilibrium as the LI seller has no incentive to deviate. If $l < \frac{1}{2}$, the buyer's optimal choice is $a(B) = B$. However, this cannot be supported in equilibrium since the LI seller can profitably deviate to $(m_{LI}(R), m_{LI}(B)) = (R, R)$.

Hence, $(m_{LI}(R), m_{LI}(B)) = (B, B), (a(R), a(B)) = (R, R)$ is an equilibrium if and only if $l \geq \frac{1}{2}$.

□

(c)

Lemma 5. *The following seller strategy:* $(m_{LI}(R), m_{LI}(B)) = (R, B)$, *cannot be supported in a PBE of the L-game where* $p = \frac{2}{3}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (R, B)$. Buyer's beliefs consistent with this strategy are: $\mu(R) = 1$, $\mu(B) = 0$. Consequently, his expected utilities from either choice are: $EU_{buyer}(a = R, m = R) = 1$, $EU_{buyer}(a = B, m = R) = 0$, $EU_{buyer}(a = R, m = B) = 0$, $EU_{buyer}(a = B, m = B) = 1$. Therefore, the buyer's best reply is: $(a(R), a(B)) = (R, B)$.

Given the buyer's best reply, the LI seller has a profitable deviation from $(m_{LI}(R), m_{LI}(B)) = (R, B)$ to $(m_{LI}(R), m_{LI}(B)) = (R, R)$ since that ensures the seller an expected payoff of 1 rather than $\frac{2}{3}$.

Therefore, the $(m_{LI}(R), m_{LI}(B)) = (R, B)$, cannot be supported in a PBE of the L-game where $p = \frac{2}{3}$.

□

(d)

Lemma 6. *The following strategies:* $(m_{LI}(R), m_{LI}(B)) = (B, R)$, $(a(R), a(B)) = (R, R)$ *are part of a PBE of the L-game where* $p = \frac{2}{3}$ *if and only if* $\frac{1}{3} \leq l \leq \frac{2}{3}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (B, R)$. Buyer's beliefs consistent with this strategy are: $\mu(R) = \frac{2-2l}{2-l}$, $\mu(B) = \frac{2l}{l+1}$.

Consequently, his expected utilities from either choice are: $EU_{buyer}(a = R, m = R) = \frac{2-2l}{2-l}$, $EU_{buyer}(a = B, m = R) = \frac{l}{2-l}$, $EU_{buyer}(a = R, m = B) = \frac{2l}{l+1}$, $EU_{buyer}(a = B, m = B) = \frac{1-l}{l+1}$.

Therefore, if $l < \frac{1}{3}$, the buyer best replies by choosing $(a(R), a(B)) = (R, B)$. However, in this case the LI seller can profitably deviate by choosing $(m_{LI}(R), m_{LI}(B)) = (R, R)$ and so, these strategies cannot represent an equilibrium.

If $l > \frac{2}{3}$, the buyer best replies by choosing $(a(R), a(B)) = (B, R)$. Again, the LI seller has a profitable deviation to $(m_{LI}(R), m_{LI}(B)) = (B, B)$ and these strategies cannot represent an equilibrium either.

If $\frac{1}{3} \leq l \leq \frac{2}{3}$, the buyer's best reply is $(a(R), a(B)) = (R, R)$. Now, the LI seller has no incentive to deviate from $(m_{LI}(R), m_{LI}(B)) = (B, R)$ and these strategies represent a PBE of the L-game where $p = \frac{2}{3}$.

□

We note that in all of the above equilibria the buyer's action towards the LI seller is R. Therefore, the LI seller's expected payoff is equal to 1 in this case. This is the same as the seller's payoff in a no-communication protocol (equivalent to the case described by Lemma 1). The buyer's payoff however can be higher than $\frac{2}{3}$ depending on the proportion of TT sellers $(1 - l)$ and the equilibrium played.

In particular, if the equilibrium play follows Lemma 3, the buyer's expected payoff is $\frac{2}{3} + \frac{1}{3}(1 - l)$. The TT seller's expected payoff in this case is equal to $\frac{2}{3}$ which is lower than the LI seller's expected payoff. In all other equilibria the players' expected payoffs are the same as in a no-communication protocol. The next proposition follows trivially.

Proposition 1. *In the L-game with* $p = \frac{2}{3}$, *if* $l < 1$, *only the buyer's best pure strategy equilibrium payoff is higher than his expected payoff in a no-communication protocol.*

Case 2. $p = \frac{1}{3}$

(a)

Lemma 7. *The following strategies:* $(m_{LI}(R), m_{LI}(B)) = (R, R)$, $(a(R), a(B)) = (R, B)$, *are part of a PBE of the L-game where* $p = \frac{1}{3}$ *if and only if* $l \leq \frac{1}{2}$. *If* $l > \frac{1}{2}$, $(m_{LI}(R), m_{LI}(B)) = (R, R)$, $(a(R), a(B)) = (B, B)$, *are part of a PBE of the L-game where* $p = \frac{1}{3}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (R, R)$. Buyer's beliefs consistent with this strategy are: $\mu(R) = \frac{1}{2l+1}$, $\mu(B) = 0$.

Consequently, his expected utilities from either choice are: $EU_{buyer}(a = R, m = R) = \frac{1}{2l+1}$, $EU_{buyer}(a = B, m = R) = \frac{2l}{2l+l}$, $EU_{buyer}(a = R, m = B) = 0$, $EU_{buyer}(a = B, m = B) = 1$.

Therefore, the buyer best replies by choosing $(a(R), a(B)) = (R, B)$ when $l \leq \frac{1}{2}$ and $(a(R), a(B)) = (B, B)$ when $l > \frac{1}{2}$. The LI seller has no incentive to deviate in either case implying that the specified strategies constitute a PBE of the L-game where $p = \frac{1}{3}$.

□

(b)

Lemma 8. *The following seller strategy:* $(m_{LI}(R), m_{LI}(B)) = (B, B)$, *cannot be supported in a PBE of the L-game where* $p = \frac{1}{3}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (B, B)$. Buyer's beliefs consistent with this strategy are: $\mu(R) = 1$, $\mu(B) = \frac{l}{2+l}$. Consequently, his expected utilities from either choice are: $EU_{buyer}(a = R, m = R) = 1$, $EU_{buyer}(a = B, m = R) = 0$, $EU_{buyer}(a = R, m = B) = \frac{l}{2+l}$, $EU_{buyer}(a = B, m = B) = \frac{2}{2+l}$.

Therefore, since $l < 1$ the buyer's best reply is: $(a(R), a(B)) = (R, B)$. The LI seller would then prefer to deviate to $(m_{LI}(R), m_{LI}(B)) = (R, R)$ proving that the specified seller strategy cannot be supported in and equilibrium of the L-game where $p = \frac{1}{3}$.

□

(c)

Lemma 9. *The following seller strategy:* $(m_{LI}(R), m_{LI}(B)) = (R, B)$, *cannot be supported in a PBE of the L-game where* $p = \frac{1}{3}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (R, B)$. Buyer's beliefs consistent with this strategy are: $\mu(R) = 1$, $\mu(B) = 0$. Consequently, his expected utilities from either choice are: $EU_{buyer}(a = R, m = R) = 1$, $EU_{buyer}(a = B, m = R) = 0$, $EU_{buyer}(a = R, m = B) = 0$, $EU_{buyer}(a = B, m = B) = 1$.

Therefore, the buyer's best reply is: $(a(R), a(B)) = (R, B)$. The LI seller would then prefer to deviate to $(m_{LI}(R), m_{LI}(B)) = (R, R)$ proving that the specified seller strategy cannot be supported in and equilibrium of the L-game where $p = \frac{1}{3}$.

□

(d)

Lemma 10. *The following strategies:* $(m_{LI}(R), m_{LI}(B)) = (B, R)$, $a(R) = a(B)$, *are part of a PBE of the L-game where* $p = \frac{1}{3}$ *if and only if* $l = \frac{1}{3}$.

Proof. Let the LI seller's strategy be $(m_{LI}(R), m_{LI}(B)) = (B, R)$. Buyer's beliefs consistent with this strategy are: $\mu(R) = \frac{1-l}{1+l}$, $\mu(B) = \frac{l}{1-l}$.

Consequently, his expected utilities from either choice are: $EU_{buyer}(a = R, m = R) = \frac{1-l}{1+l}$, $EU_{buyer}(a = B, m = R) = \frac{2l}{1+l}$, $EU_{buyer}(a = R, m = B) = \frac{l}{1-l}$, $EU_{buyer}(a = B, m = B) = \frac{1-2l}{1-l}$.

Therefore, if $l < \frac{1}{3}$, the buyer best replies by choosing $(a(R), a(B)) = (R, B)$. However, in this case the LI seller can profitably deviate to $(m_{LI}(R), m_{LI}(B)) = (R, R)$, so the specified strategies cannot constitute a PBE of the game.

If $l > \frac{1}{3}$, the buyer best replies by choosing $(a(R), a(B)) = (B, R)$. In this case the LI seller can profitably deviate to $(m_{LI}(R), m_{LI}(B)) = (B, B)$ so the specified strategies cannot constitute a PBE of the game.

If $l = \frac{1}{3}$, the buyer is indifferent between his two actions. Therefore, both $(a(R), a(B)) = (R, R)$ and $(a(R), a(B)) = (B, B)$ represent best replies to the seller's strategy. In this case, the LI seller cannot profitably deviate and the specified strategies constitute a PBE of the L-game where $p = \frac{1}{3}$.

□

In this case we note that all players can gain from communication but this again depends on the equilibrium played and on the value of l. Specifically, if the LI seller follows the strategy described in Lemma 7 and $l \leq \frac{1}{2}$, his payoff will equal 1 which is greater than his payoff in a no-communication protocol (equivalent to the one implied by Lemma 2). The TT seller's expected payoff is

$\frac{1}{3}$ which is also greater than if communication were not possible. The buyer's expected payoff is equal to $\frac{2}{3}(1-l)+\frac{1}{3}$ which is greater than $\frac{2}{3}$ given that $l \leq \frac{1}{2}$.

In all other cases, the buyer's expected equilibrium payoff is either equal to $\frac{2}{3}$ (his expected payoff in a no-communication protocol) or lower (i.e. $\frac{1}{3}$, case (d) when $(a(R), a(B)) = (R, R)$). Proposition 2 follows from this analysis.

Proposition 2. *In the L-game with $p = \frac{1}{3}$, if $l < 1$, all players' best pure strategy equilibrium payoffs are higher than their expected payoffs in a no-communication protocol.*

4.3.3.2 *Sucker cost*

We now introduce non-standard preferences for the buyers. Specifically, we assume that the buyer can incur a sucker cost whenever he chooses $a(R) = R$ after a deceptive message ($m = R$ when $s = B$). When the buyer takes an action, he does not know whether the message is deceptive or not but has an expectation about this. The higher this expectation, the more likely he is to incur the sucker cost.

Again, we restrict our analysis to only two types of buyers: those with a high enough cost which we will call "skeptics" (SK) and those whose cost is zero which we will call "non-skeptics" (NSK). Specifically, SK buyers always choose $(a_{SK}(R), a_{SK}(B)) = (B, B)$, whereas NSK buyers best reply to each message. Let $p_k \in [0,1]$ be the proportion of skeptics in the population of buyers. This, in essence, represents the probability that the buyer is a skeptic. We will call this game "the S-game". Next, we analyze the equilibria of this game given each parametrization.

Case 1. $p = \frac{2}{3}$

Lemma 11. *In every PBE of the S-game where $p = \frac{2}{3}$ and $p_k < 1$, on the equilibrium path, $(a_{NSK}(R), a_{NSK}(B)) = (R, R)$ while $(a_{SK}(R), a_{SK}(B)) = (B, B)$.*

Proof. Since the SK buyer's action is fixed, the only thing we need to consider when asking whether the NSK buyer would like to deviate from the equilibrium described by Lemma 1 is whether the proportion of SK buyers influences the optimal action for the NSK buyer. This would happen only if the seller would be incentivized to transmit more persuasive information in equilibrium the higher p_k is.

Suppose $p_k = 1$. In this case, the seller is indifferent between telling the truth or any other strategy. However, as long as $p_k < 1$, the seller would strictly prefer to send his ex ante optimal message (non-persuasive; see Lemma

1) irrespective of the state, as this makes it optimal for the NSK buyer to choose his ex ante optimal strategy $(a_{NSK}(R), a_{NSK}(B)) = (R, R)$. This way, the seller maximizes his expected payoff.

□

Case 2. $p = \frac{1}{3}$

Nothing changes in the equilibrium predictions compared to the original game. This is because the equilibrium when $p = \frac{1}{3}$ is that $a(R) = a(B) = B$ and so the NSK buyer and the SK one behave the same.

Based on the previous results we formulate the following prediction:

Proposition 3. *In the S-game where $p = \frac{2}{3}$, communication leads to material losses for the seller and the SK buyer if $p_k > 0$. Communication has no effect when $p = \frac{1}{3}$.*

Proof. The proof for why communication has no effect in the $p = \frac{1}{3}$ case of the S-game is trivial and it follows from Lemma 2, and the fact that behavior doesn't change when we account for k in this setting. To see why communication leads to losses when $p = \frac{2}{3}$, consider the players' expected payoffs in the S-game:

$$\pi_{SK_{buyer}} = \frac{1}{3}$$

$$\pi_{seller} = 1 \cdot (1 - p_k) + 0 \cdot p_k = 1 - p_k$$

Recall that the buyer's expected payoff in a no-communication protocol is $\frac{2}{3}$, while the seller's is 1. Therefore, if $p_k > 0$, both the SK buyer's and the seller's expected payoffs in the S-game are smaller than their expected payoffs in a no-communication protocol.

□

Proposition 3 implies that as long as skeptics exist, both the average buyer and the seller should lose from communication in the $p = \frac{2}{3}$ case, but not in the $p = \frac{1}{3}$ one.

4.3.4 *Hypotheses*

In this section we formulate hypotheses regarding behavior that draw from the previous theoretical analysis. We note to which type of assumption regarding players' preferences each hypothesis refers to.

Standard preferences:

The following two hypotheses are based on the game in which players have selfish, rational preferences.

Hypothesis 1. *Buyers will always choose $a = R$ when $p = \frac{2}{3}$ and $a = B$ when $p = \frac{1}{3}$, irrespective of whether communication is present.*

Hypothesis 2. *Communication has no effect on material payoffs irrespective of p.*

The second hypothesis is an implication of the fact that buyers' actions do not depend on messages, and so their presence does not affect final outcomes.

Non-standard preferences:

The following hypotheses refer to the L-game only, when only sellers have non-standard preferences.

Hypothesis 3. *When $p = \frac{2}{3}$, and a positive fraction of sellers report blue states truthfully, communication improves buyers' average material payoff and decreases the sellers' one. Communication has no effect on outcomes in this case when no truth-telling is observed.*

Hypothesis 3 assumes that the most likely equilibrium strategies in the $p = \frac{2}{3}$ case with truth-telling sellers is the one described by Lemma 3.

Hypothesis 4. *When $p = \frac{1}{3}$, and more than half of sellers report blue states truthfully, communication improves the average material payoffs of both players. Communication has no effect on outcomes in this case when less than half of sellers report blue states truthfully.*

In formulating hypothesis 4 we acknowledge the fact that the equilibrium strategies depicted by Lemma 10 are much less intuitive than those described by Lemma 7.

Next, we account for the existence of skeptical buyers.

Hypothesis 5. *When $p = \frac{2}{3}$, the share of buyers buying the blue product after a message indicating that the state is red is higher than the share of buyers choosing to buy the blue product when messages are not allowed.*

Hypothesis 6. *When $p = \frac{2}{3}$, communication reduces the average payoffs of both players.*

Finally, given the hypothesized (negative) effect of communication in the $p = \frac{2}{3}$ case and its counterpart in the $p = \frac{1}{3}$ case we can formulate a general hypothesis regarding the average effect of communication.

Hypothesis 7. *Communication has an absolute effect on outcomes.*

4.4 EXPERIMENTAL DESIGN

We run an experiment designed to test our hypotheses. This section describes the details of the design and the procedures. The experiment was conducted at the Centre for Decision Research and Experimental Economics (CeDEx) laboratory located at the University of Nottingham. The decision-making environment was computerized and the software was programmed in z-Tree (Fischbacher, 2007). There were 464 participants (out of which 270 were females) recruited from a university-wide pool of undergraduate and graduate students using ORSEE (Greiner, 2015). Participants were randomly assigned to one of the 16 sessions. Each session involved between 26 and 30 participants and lasted on average 20 minutes. All subjects received a £3 participation fee plus their earnings from playing the game. The game earnings amounted to £2 on average. In the experiment, payoffs were described in points (30 or 0) which were then exchanged to pounds according to the following exchange rate: 1 point=£0.1.

We used a 2x2 design varying the following two dimensions: the possibility of sellers to send messages and p. Table 4.2 presents a summary of the 4 treatments.

Table 4.2: Treatments summary

Name	Messages allowed?	Value of p
CM_23	YES	$\frac{2}{3}$
No_CM_23	NO	$\frac{2}{3}$
CM_13	YES	$\frac{1}{3}$
No_CM_13	NO	$\frac{1}{3}$

There were 4 sessions for each treatment. Upon arriving at the lab, participants were randomly assigned a number that would determine their seat in the lab. At their computer terminals, subjects would find a sheet of instructions[3] which were read out loud by the experimenter once everyone was seated. After the instructions were read the experiment began with a set of hypothetical control questions regarding subjects' understanding of the payoff matrix. The experiment then proceeded with the game described in Section 3.1. Participants played the game only once. Finally, they were asked to fill in a short questionnaire after which they were paid in private and in cash.

The actions of the buyer were described in terms of buying either the red or the blue product. Our communication protocol was computerized, free-form

3 See Appendix C.1 for specific wording in the CM_23 and No_CM_23 treatments. The instructions for the _13 treatments are identical except for the composition of the urn which has 2 blue balls and only 1 red ball.

and one-way in that only the seller could send a message. The seller's screen is depicted in figure 4.1. We used free form messages for comparability with Ert et al. (2014) but also to create an environment more likely to promote skepticism.

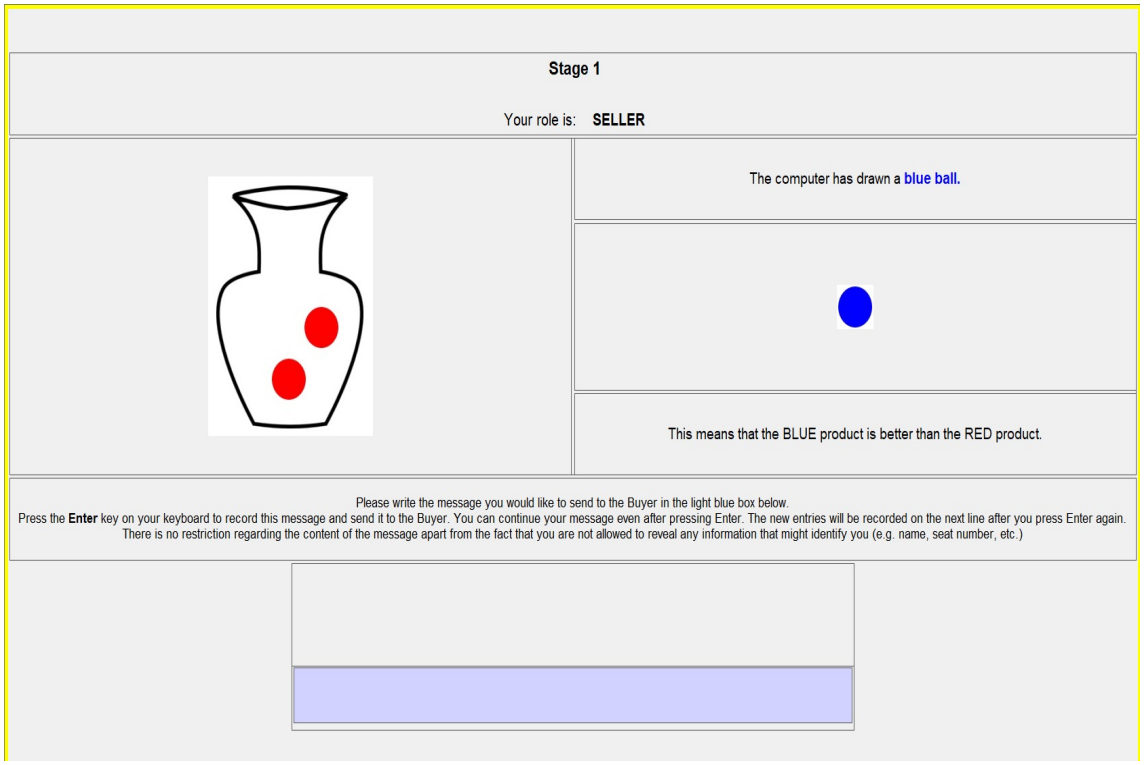


Figure 4.1: Communication interface

The state of the world was randomly determined for each subject at the start of the experiment. The exact distribution of states in each treatment is presented in Table 4.3.

Table 4.3: No. of observed red and blue states across treatments

Treatment	$s = red$	$s = blue$
CM_23	37	21
No_CM_23	42	16
CM_13	22	38
No_CM_13	19	37

4.5 RESULTS

In this section we analyze the behavior employed by buyers and seller across treatments and compare it with our hypotheses. We then investigate the effect of this behavior on payoffs. Throughout the analysis we will compare the

observed outcomes both between treatments as well as with the standard equilibrium benchmark. Since players played the game only once and there was no interaction across the buyer-seller pairs, each subject's decision represents an independent observation.

4.5.1 *Buyer behavior*

Figure 4.2 depicts how often buyers choose their ex ante optimal strategy across treatments. Recall that when $p = \frac{2}{3}$, the ex ante optimal action for the buyer is to buy the red product, whereas when $p = \frac{1}{3}$, it is to buy the blue product.

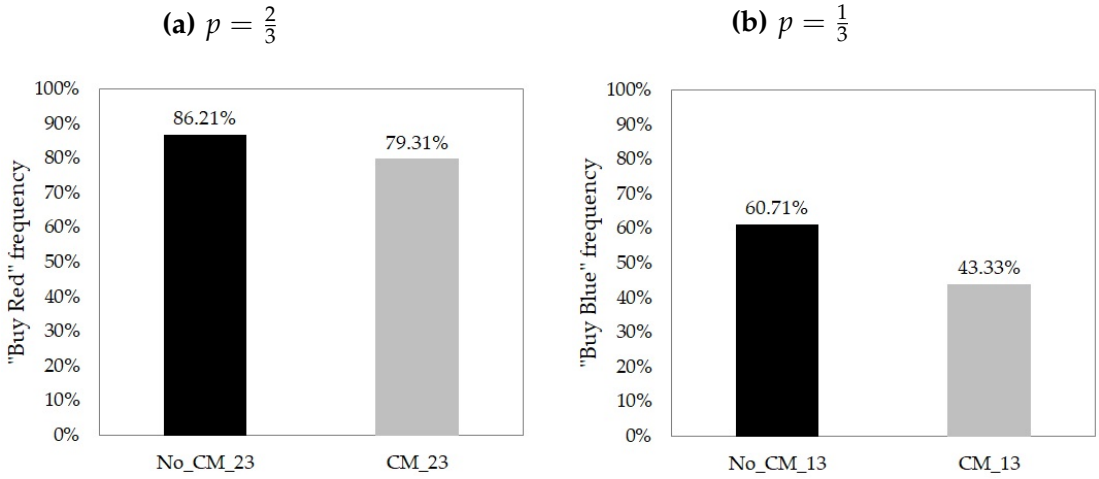


Figure 4.2: Frequency with which buyers choose the ex ante optimal action

First, we notice that communication makes buyers choose the ex ante optimal action less often. This difference is (weakly) significant in the $p = \frac{1}{3}$ case ($p - value = 0.064$, Barnard's two-sided test[4]). To better understand this finding we look at how buyers' choice of action depends on the message received. Table 4.4 reports the distribution of buyer's actions conditional on the message received.

Table 4.4: Distribution of buyer actions conditional on the message received

(a) CM_23

	$a = red$	$a = blue$
$m = red$	43	6
$m = blue$	1	4

(b) CM_13

	$a = red$	$a = blue$
$m = red$	31	11
$m = blue$	1	6

4 We also perform the Fisher test and Chi-squared test by Monte Carlo simulation for all results in this section and the findings are qualitatively the same.

We observe a significantly high proportion of buyers who choose $a = red$ when $m = red$ in both treatments. If we focus on the cases where $m = red$, the effect of communication is even stronger when $p = \frac{1}{3}$, the frequency of $a = red$ being 73.8%, which is significantly higher than 39.3% ($p - value = 0.001$, Barnard's two-sided test).

Result 1. *Buyers choose the ex ante optimal action less often in the CM_ treatments than in the No_CM_ ones. This difference is significant only in the $p = \frac{1}{3}$ case.*

The corollary of Result 1 is that communication is persuasive in the CM_13 treatment as it makes buyers take the action preferred by the seller more often. By analyzing buyers' behavior conditional on the messages received we can also find out whether communication increases buyers' skepticism in the $p = \frac{2}{3}$ case as hypothesized. We observe that the percentage of buyers choosing to buy the blue product after a red message in the CM_23 treatment is equal to 12.24%. This is significantly different from 0 ($p - value < 0.010$, one-sided binomial test). Is this evidence for sucker aversion? Recall that the validity of this claim depends on skepticism being observed only in the presence of messages. Therefore, our model implies that in the No_CM_23 treatment, this frequency would be equal to 0. Nonetheless, we find this is not the case as the frequency with which buyers buy the blue product in the no communication treatment is 13.8%. Therefore, the reason why buyers buy the blue product after a red message in the _23 treatments is not because of sucker aversion.

Result 2. *Communication does not increase skepticism when $p = \frac{2}{3}$ and therefore, the observed deviation from the ex ante optimal strategy is not due to sucker aversion.*

Result 2 goes against Hypothesis 5. Focusing on the No_CM_ treatments, an interesting finding arises: buyers choose their ex ante optimal action significantly less often than equilibrium predicts ($p - value = 0.004$ in No_CM_23 and $p - value < 0.001$ in No_CM_13, Barnard's two-sided test).

Result 3. *Buyers' behavior is significantly different than standard equilibrium predictions in the No_CM_ treatments.*

When $p = \frac{2}{3}$, the observed deviations could be due to noise as their frequency is not particularly large (13.8%). On the other hand, when $p = \frac{1}{3}$ the size of this frequency is significant (39.3%). A potential explanation is that buyers have concerns for efficiency. This could drive them to buy the red product more often than they should as this action always leads to a payoff of 1 to the seller. Taken together, this evidence rules out Hypothesis 1.

4.5.2 *Seller behavior*

This section looks at sellers' message strategy in the CM_ treatments as only there the sellers are active players. Since messages are free-form, to simplify the analysis we classify them into well-defined categories. We partition the free-form message space in three types and this accommodates the meanings of all observed messages. Table 4.6 presents the identified types and an example for each category.

Table 4.6: Message types

Type	Example
red	"Hey, I drew the red ball - the red product is better than blue!"
blue	"blue ball was drawn so pick the blue product"
inconclusive	"This product is good quality compared to the other in the urn."

Table 4.7 presents the corresponding frequencies for each treatment.

Table 4.7: Message frequencies

Treatment	'red'	'blue'	'inconclusive'
$p = \frac{2}{3}$	85%	9%	6%
$p = \frac{1}{3}$	70%	12%	18%

First, we note the use of an unexpected type of message, the inconclusive one, which does not make a claim about the color of the state. Its use could be explained if inconclusive messages are an indirect form of lying (evasive lying). Using such a message when the state is blue could give rise to lower costs than direct lying (Khalmetski et al., 2017). If this is the case, we should observe their use only when the state is blue. Table 4.8 presents how messages are distributed depending on the state of the world.

Table 4.8: Message distribution across states (cells depict overall counts)

	CM_23			CM_13		
State	'red'	'blue'	'inconclusive'	'red'	'blue'	'inconclusive'
red	33	0	4	20	0	2
blue	16	5	0	22	7	9

From table 4.8, we observe that no seller reports red states as blue. This is an indication that messages are used with their literal meanings. Second, we

see that inconclusive messages are used even when there is no incentive to lie (i.e. when the state is red). Therefore, we cannot say that these messages are a form of evasive lying. The percentage of their use when the state is red is lower than when it is blue. Nonetheless, for the rest of the analysis we will leave out the observations where an inconclusive message was sent.

Next, we note that the majority of sellers who observe a blue state report it as red. However, both in the CM_23 treatment and in the CM_13 one, we observe a significantly positive proportion of sellers who report blue states truthfully: 23.8% in CM_23 and 18.4% in CM_13 ($p - value = 0.010$ for CM_23 and $p - value = 0.003$ for CM_13, Barnard's one-sided test). This represents empirical support for the assumption that some sellers are lying averse.

Result 4. *A significant proportion of sellers (20%) report blue states truthfully.*

However, this proportion is rather small compared to typical results regarding truth-telling propensity in individual-decision making environments. Abeler et al. (2016) estimate an average amount of truth-telling at around 78%. An explanation for why the observed lying frequency in our sample is so small could be the fact that strategic environments have been found to crowd out lying aversion (Cabrales et al., 2016; Minozzi and Woon, 2013). The proportion of truth-tellers in our sample is quite similar to the 22% one observed by Sánchez-Pagés and Vorsatz (2007) in a sender-receiver game. Furthermore, we find that the likelihood of a blue state being reported as red is 58% in CM_13 and 76% in CM_23. The difference between these frequencies is only weakly statistically significant ($p - value = 0.094$, Barnard's two-sided test). Nonetheless, it is an indication that sellers perceive that there are more gains from lying in the CM_13 treatment than in the CM_23 one. This is in line with the recent findings of Gneezy et al. (2018).

Next, we examine the persuasiveness of sellers' messages as a function of the true, underlying state. To do so, we look at how often the buyers choose to buy the red product after observing a message suggesting that the state is red, conditional on the true state. From figure 4.3 we note that truthful sellers are not more persuasive than untruthful ones as there is no significant difference between the frequencies of buying the red product after a truthful versus an untruthful message.

4.5.3 *Average payoffs*

In this section we report the effect of communication on players' average payoffs. Recall that if players have standard preferences, communication should have no effect on payoffs (Hypothesis 2). In the previous sections we have

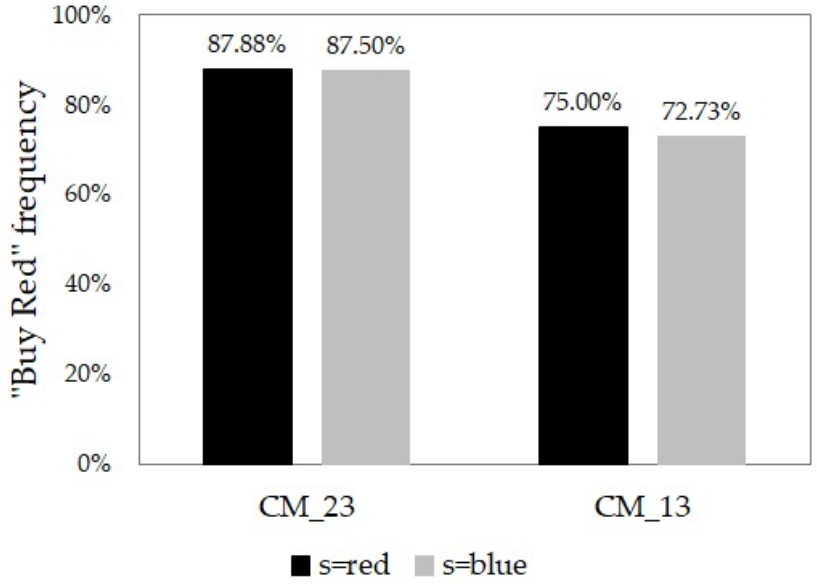


Figure 4.3: Frequency with which buyers choose $a = red$ after $m = red$ conditional on the underlying state

shown that players do not behave according to standard theoretical predictions. Nevertheless, the deviations are not high enough to render significant payoff differences according to our theoretical analysis from section 4.3.3. Table 4.9 presents the average payoffs comparison.

Table 4.9: Average payoffs across treatments

	$p = \frac{2}{3}$			$p = \frac{1}{3}$		
	CM_23	No_CM_23	p-value	CM_13	No_CM_13	p-value
Seller	0.793	0.862	0.528	0.567	0.393	0.063
Buyer	0.638	0.724	0.528	0.600	0.554	0.709

Note: p-values are based on a two-sided Barnard's test. Payoffs have been normalized to values of 1 and 0.

In line with the behavior analysis, we notice no significant payoff differences in the case of $p = \frac{2}{3}$, i.e. when interests are more likely aligned. However, when interests are more likely divergent, the seller's payoff is (weakly) significantly higher when messages are present. This is reflective of the higher proportion of buyers buying the red product in this treatment compared to the one where messages are absent. Result 5 follows:

Result 5. *When $p = \frac{2}{3}$, communication has no significant effect on players' average payoffs. When $p = \frac{1}{3}$, it significantly increases seller's average payoff and does not significantly impact buyer's average payoff.*

Result 5 contradicts Hypothesis 2. Furthermore, it goes against Hypothesis 3 since although we observe a positive share of truth-tellers, payoffs remain

unaffected (at least from a statistical perspective). The result also contradicts Hypothesis 5 as when less than half of the sellers are lying averse, communication should have no effect in the CM_13 treatment. However, it supports Hypothesis 7 according to which there is a "difference-in-differences" effect of communication.

To make sure that Result 5 is not due to differences in the underlying distribution of states across treatments, we separate observations between those where a red ball was drawn and those where a blue ball was drawn. Table 4.10 presents this information.

Table 4.10: Average payoffs across treatments conditional on the underlying state

		CM_23	No_CM_23	p-value	CM_13	No_CM_13	p-value
s=red	Seller	0.838	0.905	0.530	0.727	0.421	0.058
	Buyer	0.838	0.905	0.530	0.727	0.421	0.058
s=blue	Seller	0.714	0.750	0.842	0.474	0.378	0.529
	Buyer	0.286	0.250	0.842	0.526	0.622	0.529

First, we note that there was no significant difference between the distribution of states in the CM_ treatments versus the No_CM treatments ($p-value = 0.528$ for $p = \frac{2}{3}$ and $p-value = 0.779$ for $p = \frac{1}{3}$, Barnard's two-sided test). Second, we find that when the state is blue, payoffs are not significantly different for either player irrespective of p. The significant positive effect of communication on the seller's payoff in the case of $p = \frac{1}{3}$ arises when the state is red ($p-value = 0.058$, Barnard's two-sided test). Hence, we can conclude that communication has a positive effect for sellers not because it enables them to deceive buyers, but rather because it allows them to better signal the cases where interests are aligned.

Table 4.11 presents the comparison of players' payoffs ("Obs.") with the standard equilibrium predictions ("Eq.").

Table 4.11: Average payoffs across treatments

	Seller's average payoff			Buyer's average payoff		
	Obs.	Eq.	p-value	Obs.	Eq.	p-value
CM_23	0.793	1	<0.001	0.638	0.638	1
CM_13	0.567	0	<0.001	0.600	0.633	0.792
No_CM_23	0.862	1	0.002	0.724	0.724	1
No_CM_13	0.393	0	<0.001	0.554	0.661	0.277

Note: p-values are based on a two-sided Barnard's test.

Result 6. *Buyers' average payoff is not significantly different than equilibrium predictions. Seller's average payoff is significantly lower than predicted when* $p = \frac{2}{3}$ *and significantly higher when* $p = \frac{1}{3}$.

Result 6 seems at odds with Result 5. To understand how these can be reconciled, we look at the distribution of buyers' errors over the state space. We define an error as the case where the buyer chooses $a = red$ when $s = blue$ or $a = blue$ when $s = red$. Table 4.12 presents the observed distribution of errors alongside the predicted one. The observed distribution of errors is significantly different than the predicted one for all treatments ($p - value < 0.010$ for all four treatments, Barnard's two-sided test).

Table 4.12: Distribution of errors

Treatment	$s = red$		$s = blue$	
	Observed	Predicted	Observed	Predicted
CM_23	6	0	15	21
CM_13	6	22	18	0
No_CM_23	4	0	12	16
No_CM_13	11	19	14	0

From Table 4.12 we note that when $p = \frac{2}{3}$, the observed total number of errors is identical with the predicted one. However, buyers make a significantly higher number of errors when $s = red$. An error in this state is disadvantageous for the seller. On the other hand, when $s = blue$, buyers make fewer errors than predicted, meaning that they choose $a = blue$ more often than they should. This again is disadvantageous for the seller, but advantageous for the buyer. This is why the sellers' average payoff is significantly lower than predicted but the buyers' is not, as the increased number of errors buyers make when $s = red$ is compensated by the reduced number of errors when $s = blue$. A similar explanation holds for the case where $p = \frac{1}{3}$ as well.

In addition, we observe that the distribution of errors across the state space is clearly skewed towards $s = blue$, irrespective of the treatment. This is an indication that buyers are not able to detect deception when communication is written and anonymous, a finding different than Belot and Van de Ven (2017) where communication was face-to-face but in line with the few experiments on deception detection with written communication (Chen and Houser, 2017).

4.6 CONCLUSION

In this chapter we investigated if and when communication affects outcomes in a sender-receiver game with sender state independent preferences. We hy-

pothesized that the aversion to being profitably deceived (sucker aversion) can take over the positive effect that lying aversion might give rise to and lead to a negative effect of communication. First, we found that communication does not harm when it was hypothesized to do so (i.e. when the likelihood of interests being aligned is relatively high). This is at odds with Ert et al. (2014). A possible reason for the discrepancy in our results is the different complexity of environments used. The game implemented in Ert et al. (2014) is considerably more complex than the one we use. It could be that due to this, buyers in their experiment exert more cognitive effort to process the information. Under a higher cognitive load, equilibrium play is perhaps less likely to occur (Allred et al., 2016).

Second, our results show that communication helps when the likelihood of interests being aligned is relatively low by increasing sellers' payoff without harming buyers' payoff. Sellers benefit from the fact that buyers trust messages more often than they ought to. At the same time, buyers are not affected by their gullibility since some sellers are lying averse, reporting blue states truthfully more often than predicted by standard theory. Nevertheless, this is not predicted by our theoretical analysis since the observed amount of truth-telling is much lower than the theoretically necessary level. An explanation is that buyers overweight the probability of interacting with a lying averse seller. This is in line with Sheremeta and Shields (2017) who find that in a setting similar to ours but where the binary states are uniformly distributed, receivers are too optimistic about senders' truthful behavior. Importantly, participants in their experiment play simultaneously the role of sender and that of receiver. Hence, beliefs about senders' behavior can be confounded with self-image concerns.

Third, we find a surprising deviation from standard equilibrium predictions in the treatments where communication is absent. In the case where interests are more likely to be aligned, buyers choose the action that harms the sellers significantly more often than predicted. This is an indication that something different than sucker aversion is at play. An explanation based on probability matching (Myers, 2014; Vulkan, 2000) is ruled out as subjects do not make repeated choices. We propose two tentative explanations, none of which can capture the entire pattern of behavior. One is that buyers have different levels of strategic sophistication and those with a lower level display a higher degree of skepticism. This though is contradicting the effect we observe in the case where interests alignment is relatively less likely. In this treatment, buyers choose the action preferred by the seller significantly more often. The second possibility is that efficiency concerns matter. This would explain the latter result, but not its former counterpart.

We note, lastly, that in the treatments where interests are more likely to be aligned, the effect of efficiency concerns goes in the opposite direction with

that of communication under sucker aversion. In the treatments where interests - alignment is less likely, these effects would push behavior in the same direction. Therefore, a possibility is that the effect of efficiency concerns and of sucker aversion cancel out in the former case, while communication enhances efficiency concerns in the latter case. Further research on the interplay between these two behavioral drivers is welcomed.

BIBLIOGRAPHY

Abeler, J., Nosenzo, D., and Raymond, C. (2016). Preferences for truth-telling. *CESifo Working Paper Series No. 6087*.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.

Allred, S., Duffy, S., and Smith, J. (2016). Cognitive load and strategic sophistication. *Journal of Economic Behavior & Organization*, 125:162–178.

Bazerman, M. H., Gibbons, R., Thompson, L., and Valley, K. L. (1998). Can negotiators outperform game theory. *Debating Rationality: Nonrational Aspects of Organizational Decision Making*, pages 79–99.

Belot, M., Bhaskar, V., and Van De Ven, J. (2012). Can observers predict trustworthiness? *Review of Economics and Statistics*, 94(1):246–259.

Belot, M. and Van de Ven, J. (2017). How private is private information? the ability to spot deception in an economic game. *Experimental Economics*, 20(1):19–43.

Bicchieri, C. and Lev-On, A. (2007). Computer-mediated communication and cooperation in social dilemmas: An experimental analysis. *Politics, Philosophy & Economics*, 6(2):139–168.

Blume, A., DeJong, D. V., Kim, Y.-G., and Sprinkle, G. B. (1998). Experimental evidence on the evolution of meaning of messages in sender-receiver games. *The American Economic Review*, 88(5):1323–1340.

Blume, A., DeJong, D. V., Kim, Y.-G., and Sprinkle, G. B. (2001). Evolution of communication with partial common interest. *Games and Economic Behavior*, 37(1):79–120.

Bochet, O., Page, T., and Putterman, L. (2006). Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior & Organization*, 60(1):11–26.

Brosig, J. (2002). Identifying cooperative behavior: some experimental results in a prisoner's dilemma game. *Journal of Economic Behavior & Organization*, 47(3):275–290.

Burton, A., Loomes, G., and Sefton, M. (2005). Communication and efficiency in coordination game experiments. *Experimental and Behavorial Economics*, pages 63–85.

Cabrales, A., Feri, F., Gottardi, P., and Melendez-Jimenez, M. (2016). Can there be a market for cheap-talk information? Some experimental evidence. *Mimeo*.

Cai, H. and Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36.

Charness, G. (2000). Self-serving cheap talk: A test of Aumann's conjecture. *Games and Economic Behavior*, 33(2):177–194.

Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6):1579–1601.

Charness, G. and Dufwenberg, M. (2011). Participation. *American Economic Review*, 101(4):1211–37.

Chen, J. and Houser, D. (2017). Promises and lies: can observers detect deception in written messages. *Experimental Economics*, 20(2):396–419.

Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, 107(2):739–771.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Darke, P. R. and Ritchie, R. J. (2007). The defensive consumer: Advertising deception, defensive processing, and distrust. *Journal of Marketing Research*, 44(1):114–127.

Dawes, R. M., McTavish, J., and Shaklee, H. (1977). Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation. *Journal of Personality and Social Psychology*, 35(1):1.

Dickhaut, J. W., McCabe, K. A., and Mukherji, A. (1995). An experimental study of strategic information transmission. *Economic Theory*, 6(3):389–403.

Ert, E., Creary, S., and Bazerman, M. (2014). Cynicism in negotiation: When communication increases buyers' skepticism. *Judgment and Decision Making*, 9:191–198.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Forsythe, R., Lundholm, R., and Rietz, T. (1999). Cheap talk, fraud, and adverse selection in financial markets: Some experimental evidence. *The Review of Financial Studies*, 12(3):481–518.

Friestad, M. and Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research*, 21(1):1–31.

Gneezy, U., Kajackaite, A., and Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, 108(2):419–53.

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

Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies*, 76(4):1359–1395.

Khalmetski, K., Rockenbach, B., and Werner, P. (2017). Evasive lying in strategic communication. *Journal of Public Economics*, 156:59–72.

Lipnowski, E. and Ravid, D. (2017). Cheap talk with fixed agenda. *Mimeo*.

Minozzi, W. and Woon, J. (2013). Lying aversion, lobbying, and context in a strategic communication experiment. *Journal of Theoretical Politics*, 25(3):309–337.

Mohlin, E. and Johannesson, M. (2008). Communication: Content or relationship? *Journal of Economic Behavior & Organization*, 65(3-4):409–419.

Myers, J. L. (2014). Probability learning and sequence learning. *Handbook of Learning and Cognitive Processes, ed. WK Estes*, pages 171–205.

Ockenfels, A. and Selten, R. (2000). An experiment on the hypothesis of involuntary truth-signalling in bargaining. *Games and Economic Behavior*, 33(1):90–116.

Roth, A. E. (1995). *Bargaining experiments*. Princeton University Press.

Sánchez-Pagés, S. and Vorsatz, M. (2007). An experimental study of truth-telling in a sender–receiver game. *Games and Economic Behavior*, 61(1):86–112.

Serra-Garcia, M., van Damme, E., and Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, 73(1):244–261.

Sheremeta, R. M. and Shields, T. W. (2017). Deception and reception: The behavior of information providers and users. *Journal of Economic Behavior &*

Organization, 137:445–456.

Valley, K., Thompson, L., Gibbons, R., and Bazerman, M. H. (2002). How communication improves efficiency in bargaining games. *Games and Economic Behavior*, 38(1):127–155.

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14(1):101–118.

5

GENERAL CONCLUSIONS

This thesis contributed to the understanding of strategic behavior in environments where one party can use non-costly communication to influence the actions of another. In all chapters we used an experimental methodology based on sender-receiver games to investigate the informativeness of theoretical predictions.

In chapter 2 we compared behavior in two settings where both verifiable and non-verifiable bi-dimensional messages about the state of the world are present. In fact, one of the contributions of this chapter is the experimental investigation of behavior when more than one type of communication is available. The setups under scrutiny differ in whether the receiver or the sender has control over the verification action. Importantly, the verification process can reveal only one dimension of the state which allows the unverifiable message to influence beliefs about the true state. Nevertheless, to achieve his best equilibrium payoff the receiver should ignore the cheap talk upon having observed the partially-informative evidence. This theoretical result does not depend on who has verification control and the receiver's best equilibrium outcomes are identical across settings (Glazer and Rubinstein, 2004, 2006). Notwithstanding, since equilibrium structures differ qualitatively between these two settings, we hypothesize that equilibrium outcomes are more likely to emerge in the case where the sender rather than the receiver controls the verification action.

Indeed, our main finding corroborates this hypothesis. Behavior is closer to the equilibrium benchmark when the sender has verification control although payoffs are not significantly different between settings. However, giving control over the verification action to the receiver rather than to the sender, promotes payoffs closer to equilibrium predictions for both players. The receiver is hurt given the equilibrium benchmark in the other setting. The peculiarity of this result stems from how the sender's deviations from equilibrium play complements the receiver's when the latter has verification control. The robustness of this result is worth exploring through a replication study. Another interesting finding is that receivers are systematically less skeptical when they are in control of their information space.

Overall, this chapter highlights the potential gains that would arise if some informational control is transferred to the receiver. In our setting senders do

not lose from this process. In more naturalistic environments though, this is perhaps rather costly given the larger variance of verification strategies. Companies and institutions might consequently find it more difficult to predict subsequent costs and use their resources efficiently. A suggestion for future research would thus be to investigate behavior when information transmission is costly in frameworks with both verifiable and unverifiable communication.

To deepen our understanding of receiver's behavior in chapter 2, the next chapter focused on the effect of unverifiable messages (cheap talk) on the interpretation of evidence (verifiable communication). To examine this, we used the setting from chapter 2 where sender's verification control is coupled with the ability to send cheap talk messages. We then compare this framework with a similar one where senders can only transmit verifiable messages. We find that receivers are influenced by cheap talk messages more than equilibrium predicts. Furthermore, these are an important driver of receiver's noisy behavior and consequently of negative welfare effects. In a follow-up treatment we reconcile the result from chapter 2 according to which giving receivers more control over their information space reduces unprofitable skepticism and the payoff gap due to cheap talk.

Although research on the effect of cheap communication is vast, this chapter examined a novel avenue providing useful insights into how individuals might weigh different sources of information in their decision process. One limitation is that we do not observe directly receivers' beliefs but rather infer them from their actions. This is arguably a "cleaner" approach as direct belief elicitation techniques often give rise to confounding results. Notwithstanding, in our setting, a more in depth analysis of receivers' beliefs could identify other factors that influence their decisions. Other-regarding preferences for example seem to be particularly relevant in strategic settings. We welcome future research in this direction.

Chapter 4 further explores the welfare implications of communication. Inspired by previous findings, we identify two non-standard preferences and show theoretically that each can give rise to a different effect. On the one hand, "sucker aversion" can lead to unprofitable skepticism (see chapter 3 and Ert et al., 2014). On the other hand, "lying aversion" can bring about gains from communication (Abeler et al., 2016). These predictions depend on whether the interests of the two players are more or less likely to be aligned. We test these predictions by comparing a cheap talk game with sender state independent preferences to a similar setting where unverifiable messages are absent. In this environment, if players have standard preferences and the receiver is not indifferent between his actions, communication should be inconsequential irrespective of the degree of interests alignment (Lipnowski and Ravid, 2017).

The main result of this chapter is that communication matters. In particular, we find that senders' payoff is positively affected when interests are more likely conflicting and not affected at all when interests are more likely aligned. Communication never influences the receiver's payoff. Even though the hypothesized mechanism (skepticism due to "sucker aversion") is not supported by our findings, this chapter highlights that non-standard preferences are important, efficiency concerns being a potential driver of the main finding. At face value, one could find it surprising that skepticism is not as prevalent in this setting as it was in the previous chapters or in Ert et al. (2014). One potential reason may be the complexity of the games. Experimental settings in chapters 2 and 3 are significantly more complex when compared to that in chapter 4. Perhaps more complex environments that are more cognitively taxing, are less likely to promote equilibrium play (Allred et al., 2016). There are several aspects that contribute to the complexity of these games such as the size of the state space and the communication protocol. Future research could try to disentangle among these in search for potential causes of our findings.

Finally, this thesis showed that accurately predicting behavior in settings of strategic communication is a tall task. Nevertheless, given enough empirical data, patterns emerge which can inform new theoretical models that are instrumental for the design of efficient markets and institutions.

BIBLIOGRAPHY

Abeler, J., Nosenzo, D., and Raymond, C. (2016). Preferences for truth-telling. *CESifo Working Paper Series No. 6087*.

Allred, S., Duffy, S., and Smith, J. (2016). Cognitive load and strategic sophistication. *Journal of Economic Behavior & Organization*, 125:162–178.

Ert, E., Creary, S., and Bazerman, M. (2014). Cynicism in negotiation: When communication increases buyers' skepticism. *Judgment and Decision Making*, 9:191–198.

Glazer, J. and Rubinstein, A. (2004). On optimal rules of persuasion. *Econometrica*, 72(6):1715–1736.

Glazer, J. and Rubinstein, A. (2006). A study in the pragmatics of persuasion. *Theoretical Economics*, 1:395–410.

Lipnowski, E. and Ravid, D. (2017). Cheap talk with fixed agenda. *Mimeo*.

A

APPENDIX FOR CHAPTER 2

A.1 EQUILIBRIUM EXAMPLES

In order to illustrate the multiplicity of equilibria, in this section we describe a few examples. First, we note that the GR equilibrium is the receiver's best equilibrium both in Sr and in Rv, for either parametrization (T=11 or T=9).

- One other equilibrium in Rv is the babbling one. In a babbling equilibrium, the message is uninformative (e.g. all sender types send each of the possible messages with equal probability, or all sender types send the same message). The receiver does not condition the checking decision on the message (e.g. the receiver always checks the same card, or checks each of the two cards with equal probability). Upon observing the true value of one of the cards, the receiver takes the optimal acceptance decision given that the other card is equally likely to be any value between 1 and 9. For T = 9, the receiver accepts if and only if the observed card is at least 4 (since in 5/9 of cases the sum is at least 9). This equilibrium gives rise to a payoff of $\frac{62}{81}$ for the receiver and $\frac{54}{81}$ for the sender. For T = 11, the receiver accepts if and only if the observed card is at least 6 (since in 5/9 of cases the sum is at least 11). This equilibrium gives rise to a payoff of $\frac{61}{81}$ for the receiver and $\frac{36}{81}$ for the sender.

- An equilibrium different from GR in SR when T=11 involves the sender revelation strategy depicted in table A.1. This strategy implies that the sender reveals the highest card for bad hands and the lower one for good hands (when the two cards are equal, one is revealed at random). Irrespective of the value of the observed card, this gives rise to a probability of the hand being good conditional on the observed value, lower than 0.5. Hence, the receiver best replies by always rejecting, which leads to a number of 46 errors. This gives rise to a payoff of $\frac{45}{81}$ for the receiver and 0 for the sender.

- An equilibrium different from GR in SR when T=9 involves the sender revelation strategy depicted in table A.2. Again, the sender reveals the highest card when the hand is bad and the lowest one when the hand is good (mixing with equal probability between the two in case of equality). The receiver best replies by always accepting, a strategy which leads to

a total number of errors equal to 28. This gives rise to a payoff of $\frac{53}{81}$ for the receiver and 1 for the sender.

Table A.1: Sender strategy for receiver's worst equilibrium when T=11

9	9	**2**	**3**	**4**	**5**	**6**	**7**	**8**	9/**9**
8	8	8	**3**	**4**	**5**	**6**	**7**	8/**8**	8
7	7	7	7	**4**	**5**	**6**	7/**7**	7	7
6	6	6	6	6	**5**	6/**6**	6	6	6
5	5	5	5	5	5/**5**	5	5	5	5
4	4	4	4	4/**4**	**5**	**6**	4	4	4
3	3	3	3/**3**	**4**	**5**	**6**	**7**	3	3
2	2	2/**2**	**3**	**4**	**5**	**6**	**7**	**8**	2
1	1/**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
orange card/**blue card**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**

Table A.2: Sender strategy for receiver's worst equilibrium when T=9

9	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	9/**9**
8	**1**	**2**	**3**	**4**	**5**	**6**	**7**	8/**8**	8
7	7	**2**	**3**	**4**	**5**	**6**	7/**7**	7	7
6	6	6	**3**	**4**	**5**	6/**6**	6	6	6
5	5	5	5	**4**	5/**5**	5	5	5	5
4	4	4	4	4/**4**	4	4	4	4	4
3	3	3	3/**3**	**4**	**5**	3	3	3	3
2	2	2/**2**	**3**	**4**	**5**	**6**	2	2	2
1	1/**1**	**2**	**3**	**4**	**5**	**6**	**7**	1	1
orange card/**blue card**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**

A.2 AVERAGE PAYOFFS PER MATCHING GROUP

Tables A.3 and A.4 present the average payoffs for each player within each matching group. These represent our independent observations.

Table A.3: Average payoffs per matching group: T=11

	Receiver			Sender		
Matching group	Rv	Sr	GR prediction	Rv	Sr	GR prediction
1	0.822	0.717	0.833	0.606	0.578	0.628
2	0.722	0.789	0.806	0.528	0.439	0.544
3	0.856	0.744	0.794	0.506	0.506	0.556
4	0.794	0.783	0.817	0.600	0.567	0.567
5	0.744	0.761	0.800	0.528	0.467	0.561
6	0.794	0.822	0.850	0.589	0.506	0.544
7	0.717	0.789	0.817	0.494	0.578	0.506
8	0.806	0.733	0.789	0.478	0.561	0.528

Table A.4: Average payoffs per matching group: T=9

	Receiver			Sender		
Matching group	Rv	Sr	GR prediction	Rv	Sr	GR prediction
1	0.860	0.853	0.900	0.693	0.687	0.627
2	0.907	0.820	0.893	0.707	0.660	0.640
3	0.872	0.800	0.889	0.706	0.644	0.667
4	0.872	0.789	0.889	0.711	0.639	0.694
5	0.867	0.839	0.906	0.706	0.633	0.678
6	0.839	0.889	0.894	0.722	0.772	0.700
7	0.907	0.787	0.860	0.660	0.673	0.707
8	0.853	0.860	0.900	0.707	0.713	0.713

A.3 RECEIVER'S BEST RESPONSE GIVEN THE EMPIRICAL DISTRIBUTION IN RV, T=9

The following table excludes the messages where either of the two reports is 8 or a 9. In these cases, it is optimal to check the higher report and accept if it is true as this strategy gives a probability of a good hand conditional on the report being true equal to 100%.

Table A.5: Calculating the best response to the empirical distribution for Rv, T=9

Message	Absolute frequency of message	Absolute frequency of good hands given message	% Good hands given highest report is true	% Good hands given lowest report is true	Expected payoff from: check highest report & accept if true[1]	Expected payoff from: check lowest report & accept if true[2]	Expected payoff from: check either report & reject[3]
(7,7)	30	23	100%	100%	0.817	0.817	0.233
(7,6)	42	37	97%	94%	0.881	0.786	0.119
(7,5)	43	41	100%	95%	1.000	0.930	0.047
(7,4)	32	24	89%	83%	0.906	0.844	0.250
(7,3)	52	28	77%	67%	0.827	0.731	0.462
(7,2)	70	23	68%	40%	0.843	0.500	0.671
(6,6)	55	38	86%	86%	0.636	0.636	0.309
(6,5)	47	34	94%	74%	0.894	0.681	0.277
(6,4)	47	29	81%	72%	0.851	0.745	0.383
(6,3)	73	30	68%	51%	0.808	0.603	0.589
(5,5)	40	16	57%	57%	0.688	0.688	0.600
(5,4)	125	37	47%	47%	0.664	0.664	0.704
				Total expected payoff	0.796	0.692	0.451

1 Computed by averaging over the instances in which the highest report is untrue and the hand is bad and those where the highest report is true and the hand is good. For each such instance the receiver earns 1 point which is then weighted by the frequency of the corresponding message.

2 Computed by averaging over the instances in which the lowest report is untrue and the hand is bad and those where the lowest report is true and the hand is good. For each such instance the receiver earns 1 point which is then weighted by the frequency of the corresponding message.

3 Computed by adding the number of bad hands and then dividing by the frequency of the corresponding message.

A.4 INSTRUCTIONS

A.4.1 *Rv, T=11*

INSTRUCTIONS

Welcome and thank you for participating in this experiment. Throughout the whole experiment you are kindly asked to remain seated and refrain from communication with the other participants. Mobile phones and other electronic devices should be switched off. If there are any questions please raise your hand and an experimenter will come to answer your questions in private.

Payment: This experiment consists of 30 rounds. In each round you can earn points. At the end of the experiment you will be paid according to your accumulated point-earnings from all rounds. You will be paid in private and in cash with **£0.50 for each point earned**. Additionally, you will receive a **participation fee of £3.**

All your decisions are anonymous, so your identity will be kept secret at all times.

At the beginning of each round you will be randomly matched with another participant (i.e. the person you are paired with will change from round to round). One of you will have the role of Person A and the other the role of Person B. Your role will be assigned at the beginning of the first round and you will keep this role for all 30 rounds.

Each round consists of 2 stages which are described below.

Stage 1*: Person A observes two cards and sends a message*

In this stage, the computer will randomly select two cards, one **orange** and one **blue**, each carrying a value between 1 and 9. Each combination of values on these 2 cards is equally probable. At this stage, only **Person A** will be able to observe these values.

The hand is "**GOOD"** if the sum of the values on the two cards is **at least 11**. The hand is "**BAD**" if the sum of the values is **10 or below**.

This is an example of a BAD hand:



After observing the randomly drawn cards, Person A will send a message to Person B of the following form:

The value of the **orange card** is: ☐

The value of the **blue card** is: ☐

where Person A fills each blank box with a number between 1 and 9.

Stage 2: *Person B selects one of the cards to observe and makes a decision*

After observing Person A's message, Person B will **select one of the two cards** (**orange** or **blue**) and the computer will reveal its value.

After observing the value on the selected card Person B will **decide between "Accept" and "Reject"**.

End of the Round

At the end of the round both Person A and Person B will receive a summary of the round including:

- The cards that were randomly dealt;
- Person A's message;
- Person B's choice regarding which card to observe;
- Person B's decision to accept or reject;
- Person A and Person B's point-earnings for the round.

How your point earnings are determined:

Person A earns 1 point if B accepts and 0 points if B rejects.

Person B earns 1 point if A has a good hand and B accepts or if A has a bad hand and B rejects. Person B earns 0 points otherwise.

This is summarised in the Table below:

	B **Accepts**	B **Rejects**
A has a **GOOD** hand	Person A receives 1 point, Person B receives 1 point	Person A receives 0 points, Person B receives 0 points
A has a **BAD** hand	Person A receives 1 point, Person B receives 0 points	Person A receives 0 points, Person B receives 1 point

Preliminary questions: Before the 30 rounds begin, you will be asked to answer a few questions regarding your understanding of the instructions. The rounds will begin only after all participants have answered these questions correctly.

Final questionnaire: After the 30 rounds, you will be asked to fill in a short questionnaire. You will then be paid your earnings in private and in cash.

A.4.2 *Sr, T=11*

INSTRUCTIONS

Welcome and thank you for participating in this experiment. Throughout the whole experiment you are kindly asked to remain seated and refrain from communication with the other participants. Mobile phones and other electronic devices should be switched off. If there are any questions please raise your hand and an experimenter will come to answer your questions in private.

Payment: This experiment consists of 30 rounds. In each round you can earn points. At the end of the experiment you will be paid according to your accumulated point-earnings from all rounds. You will be paid in private and in cash with **£0.50 for each point earned**. Additionally, you will receive a **participation fee of £3.**

All your decisions are anonymous, so your identity will be kept secret at all times.

At the beginning of each round you will be randomly matched with another participant (i.e. the person you are paired with will change from round to round). One of you will have the role of Person A and the other the role of Person B. Your role will be assigned at the beginning of the first round and you will keep this role for all 30 rounds.

Each round consists of 3 stages which are described below.

Stage 1*: Person A observes two cards and sends a message*

In this stage, the computer will randomly select two cards, one **orange** and one **blue**, each carrying a value between 1 and 9. Each combination of values on these 2 cards is equally probable. At this stage, only **Person A** will be able to observe these values.

The hand is "**GOOD"** if the sum of the values on the two cards is **at least 11**. The hand is "**BAD**" if the sum of the values is **10 or below**.

This is an example of a BAD hand:



After observing the randomly drawn cards, Person A will send a message to Person B of the following form:

The value of the **orange card** is: ☐

The value of the **blue card** is: ☐

where Person A fills each blank box with a number between 1 and 9.

Stage 2: *Person A selects one of the cards for Person B to observe*

After Person B observes Person A's message, Person A will **select one of the two cards** (**orange** or **blue**) for Person B to observe its value in the next stage.

Stage 3: *Person B observes the value of the card and makes a decision*

After observing the value of the selected card, Person B will **decide between "Accept" and "Reject"**.

End of the Round

At the end of the round both Person A and Person B will receive a summary of the round including:

- The cards that were randomly dealt;
- Person A's message;
- Person A's choice regarding which card to be observed by Person B;
- Person B's decision to accept or reject;
- Person A and Person B's point-earnings for the round.

How your point earnings are determined:

Person A earns 1 point if B accepts and 0 points if B rejects.

Person B earns 1 point if A has a good hand and B accepts or if A has a bad hand and B rejects. Person B earns 0 points otherwise.

This is summarised in the Table below:

	B **Accepts**	B **Rejects**
A has a **GOOD** hand	Person A receives 1 point, Person B receives 1 point	Person A receives 0 points, Person B receives 0 points
A has a **BAD** hand	Person A receives 1 point, Person B receives 0 points	Person A receives 0 points, Person B receives 1 point

Preliminary questions: Before the 30 rounds begin, you will be asked to answer a few questions regarding your understanding of the instructions. The rounds will begin only after all participants have answered these questions correctly.

Final questionnaire: After the 30 rounds, you will be asked to fill in a short questionnaire. You will then be paid your earnings in private and in cash.

B

APPENDIX FOR CHAPTER 3

B.1 ITERATED ELIMINATION OF WEAKLY DOMINATED STRATEGIES IN NO_CM TREATMENTS

Note that the game used in the No_CM treatments from chapter 3 is the same as the Sr game in chapter 2 where messages are not allowed.

Case 1: T=9

If the receiver observes a card with a value of 8 or 9, the hand is certain to be a good hand; hence, any strategies of the receiver that allow for rejecting a hand after observing a value of 8 or 9 of either color are weakly dominated. Given that the receiver must accept cards with values 8 or 9, any strategies where the sender may display some other value even though an 8 or a 9 is available are weakly dominated.

The next step is to eliminate all strategies where the receiver accepts a hand after observing a value of 1. A value of 1 is not automatically a bad hand (the other card could be 8 or 9) but, because of the previous step, if the other card was 8 or 9 the sender would have displayed it. Hence, because of the previous stage of elimination, a value of 1 is sure to be a bad hand and any receiver strategies where a 1 is accepted with positive probability are eliminated.

Given that the receiver rejects cards with a value of 1, it becomes weakly dominated for the sender to show a 1 if the other card is not a 1. In particular, if the hand is (1,2) or (2,1), the 2 is always shown. We can now eliminate any strategies where the receiver accepts with positive probability after observing 2. Of the three good hands that contain a 2, the ones where the other card is 8 or 9 are ruled out by the sender's strategy. Of the bad hands, the one where the other card is a 1 always results in 2 being observed. Hence, the probability that the hand is good conditional on a value of 2 being observed cannot exceed 0.5 given the surviving strategies of the sender (the best-case scenario occurs if the 2 is always shown when the other card has a value of either 7 or 1, and never otherwise). Since the receiver is at best indifferent, strategies where the receiver accepts a value of 2 are weakly dominated and can be eliminated.

Given that a 2 is rejected for sure, it follows that the sender will display the other card if it is higher than 2. It then becomes weakly dominated for the receiver to reject 7: the worst that could happen given the surviving strategies for the sender is that the other card is equally likely to be a 1 and a 2; the receiver is indifferent in this case and would strictly prefer to accept in all other cases. For similar reasons, it becomes weakly dominated for the receiver to accept 3: given the surviving strategies of the sender, there are two bad hands that are showing 3 for sure, and only two good hands that might be showing a 3.

Given that the receiver accepts 7 and rejects 3, the sender will display a 7 (unless the other card is 8 or 9) and hide a 3 (unless the other card is 1 or 2). In the next step, the receiver rejects 4 (which is now unambiguously more likely to be a bad hand) and 5 (which is at best equally likely to be a good or a bad hand). The sender then hides these cards (unless the other card is lower; in cases (4,5) and (5,4) any of the two cards can be shown). After this is known, we can eliminate any strategies for the receiver in which 6 is rejected.

Both players are indifferent between all action profiles that survive this procedure, hence the game is dominance solvable. There is still some indeterminacy in the strategy of the sender (which card to show when both cards have the same value, and also when the values are 4 and 5 or 8 and 9). In all the surviving strategy profiles, senders who have a 6+ card are accepted for sure, and senders who do not are rejected for sure.

Case 2: T=11

The case T = 11 is analogous. The first step in this case is to eliminate any strategies in which the receiver accepts with positive probability after observing a 1 (since this is a guaranteed bad hand). The sender can then eliminate all the strategies where the highest card is not shown for sure even though the other card is a 1. The next step is to notice that, when a 2 of either color is observed, the probability that the hand is good cannot exceed 50% (when the other card is a 1, the 2 is always shown, and there is only one good hand that might be displaying a 2, which occurs when the other card is a 9); it is then weakly dominated to accept a 2. The sender will then display the highest card if the lowest card is a 2. The receiver then accepts 9 (since there are two hands that are displaying the 9 for sure, of which one is a good hand and the other a bad hand, and all other hands are good; hence, in the worst possible case the receiver would be indifferent between accepting and rejecting) and rejects 3 (there are only two good hands compatible with a value of 3, and two bad hands that are displaying the 3 for sure; hence, in the best possible case the receiver would be indifferent). The sender can then eliminate all strategies where a 9 is not displayed, or where a 3 is displayed even though the other card is higher. Taking this into account, the receiver can

eliminate all strategies where 4 or 5 is accepted, and the sender can eliminate all strategies where a 4 or a 5 is displayed even though the other card is 6+. It then follows that the receiver must accept 8 (and the sender displays 8 whenever the other card is lower). The receiver then rejects 6, and the sender displays the higher card whenever the lower card is a 6. The receiver then accepts 7.

As for the case T = 9, both players are indifferent between all action profiles that survive this procedure, hence the game is dominance solvable. There is still some indeterminacy in the strategy of the sender (which may show any of the two cards when both cards have the same value, and also when the values are 4 and 5). In all the surviving strategy profiles, senders who have a 7+ card are accepted for sure, and senders who do not are rejected for sure.

B.2 INSTRUCTIONS (NO_CM_9)

INSTRUCTIONS

Welcome and thank you for participating in this experiment. Throughout the whole experiment you are kindly asked to remain seated and refrain from communication with the other participants. Mobile phones and other electronic devices should be switched off. If there are any questions please raise your hand and an experimenter will come to answer your questions in private.

Payment: This experiment consists of 30 rounds. In each round you can earn points. At the end of the experiment you will be paid according to your accumulated point-earnings from all rounds. You will be paid in private and in cash with **£0.50 for each point earned**. Additionally, you will receive a **participation fee of £3.**

All your decisions are anonymous, so your identity will be kept secret at all times.

At the beginning of each round you will be randomly matched with another participant (i.e. the person you are paired with will change from round to round). One of you will have the role of Person A and the other the role of Person B. Your role will be assigned at the beginning of the first round and you will keep this role for all 30 rounds.

Each round consists of 3 stages which are described below.

Stage 1*: Person A observes two cards.*

In this stage, the computer will randomly select two cards, one **orange** and one **blue**, each carrying a value between 1 and 9. Each combination of values on these 2 cards is equally probable. At this stage, only **Person A** will be able to observe these values.

The hand is "**GOOD"** if the sum of the values on the two cards is **at least 9**. The hand is "**BAD**" if the sum of the values is **8 or below**.

This is an example of a BAD hand:



Stage 2: *Person A selects one of the cards for Person B to observe.*

After observing the randomly drawn cards Person A will **select one of the two cards** (**orange** or **blue**) for Person B to observe its value in the next stage.

Stage 3: *Person B observes the value of the card and makes a decision.*

After observing the value on the selected card Person B will **decide between "Accept" and "Reject"**.

End of the Round

At the end of the round both Person A and Person B will receive a summary of the round including:

- The cards that were randomly dealt;
- Person A's choice regarding which card to be observed by Person B;
- Person B's decision to accept or reject;
- Person A and Person B's point-earnings for the round.

How your point earnings are determined:

Person A earns 1 point if B accepts and 0 points if B rejects.

Person B earns 1 point if A has a good hand and B accepts or if A has a bad hand and B rejects. Person B earns 0 points otherwise.

This is summarised in the Table below:

	B **Accepts**	B **Rejects**
A has a **GOOD** hand	Person A receives 1 point, Person B receives 1 point	Person A receives 0 points, Person B receives 0 points
A has a **BAD** hand	Person A receives 1 point, Person B receives 0 points	Person A receives 0 points, Person B receives 1 point

Preliminary questions: Before the 30 rounds begin, you will be asked to answer a few questions regarding your understanding of the instructions. The rounds will begin only after all participants have answered these questions correctly.

Final questionnaire: After the 30 rounds, you will be asked to fill in a short questionnaire. You will then be paid your earnings in private and in cash.

C

APPENDIX FOR CHAPTER 4

C.1 INSTRUCTIONS

C.1.1 *CM_23*

INSTRUCTIONS

Welcome and thank you for participating in this experiment. Throughout the whole experiment you must remain seated and mobile phones and other electronic devices must be switched off. If there are any questions please raise your hand and an experimenter will come to answer your questions in private.

Payment: In this experiment you can earn points. At the end of the experiment you will be paid **£3 as a participation fee** plus **£0.10 for each point earned**. You will be paid in private and in cash**.**

All your decisions are anonymous, so your identity will be kept secret at all times.

Brief description of the task:

At the beginning of the experiment you will be randomly matched with another participant. One of you will have the role of "Buyer" and the other the role of "Seller".

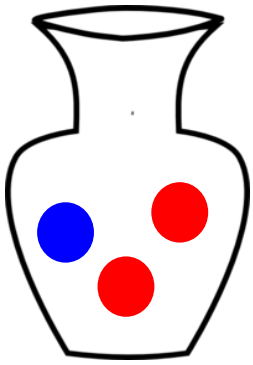
There are two products: a Red product and a Blue product. The Buyer must buy one of these products. The Seller knows the qualities of the two products but the Buyer does not. The final payoff to both participants depends on the buying choice and the quality of the chosen product. A more detailed description of the task is given below.

Detailed description of the task:

The task consists of 2 stages.

Stage 1: *The Seller is informed about the qualities of the two products and sends a message.*

In this stage, the computer will randomly draw a ball from a computerized urn containing 3 balls: 2 red and 1 blue. The urn is depicted below:



The Seller will observe the colour of the drawn ball, but the Buyer will not.

- If the colour of the drawn ball is **red**, the **Red product** is **better than** the **Blue product.** The Red product is worth 30 points to the Buyer and the Blue product is worth nothing to the Buyer.
- If the colour of the drawn ball is **blue**, the **Blue product** is **better than** the **Red product.** The Blue product is worth 30 points to the Buyer and the Red product is worth nothing to the Buyer.

After observing the qualities of the two products, the Seller will send a message to the Buyer, by typing in a chat box. There is no restriction regarding the content of the message, except that Sellers are not allowed to reveal any information that might identify them (e.g. name, seat number, etc.)

Stage 2: *The Buyer observes the message and decides whether to buy the Red product or the Blue product.*

After the Seller observes the qualities of the products and after observing the message from the Seller, the Buyer will decide between buying either the Red product or the Blue product.

How your earnings are determined:

The Buyer will earn 30 points if they buy the better product. The Seller will earn 30 points if the Buyer buys the Red product. Thus, the point-earnings will be determined as follows.

The Buyer

- earns 30 points if a Red ball is drawn and she buys the Red product;
- earns nothing if a Red ball is drawn and she buys the Blue product;
- earns 30 points if a Blue ball was drawn and she buys the Blue product;
- earns nothing if a Blue ball was drawn and she buys the Red product.

The Seller

- earns 30 points if the Buyer buys the Red product;
- earns nothing if the Buyer buys the Blue product.

This is summarised in the Table below:

	Buy the Red product	**Buy the Blue product**
Red ball drawn	The Seller earns 30 points, The Buyer earns 30 points	The Seller earns 0 points, The Buyer earns 0 points
Blue ball drawn	The Seller earns 30 points, The Buyer earns 0 points	The Seller earns 0 points, The Buyer earns 30 points

Preliminary questions: Before the experiment begins, you will be asked to answer a few questions regarding your understanding of the instructions. The experiment will begin only after all participants have answered these questions correctly.

Final questionnaire: After everyone makes their decisions, you will be asked to fill in a short questionnaire. You will then be paid your earnings in private and in cash. You will receive a £3 participation fee plus £0.10 for each point that you earn.

C.1.2 *No_CM_23*

INSTRUCTIONS

Welcome and thank you for participating in this experiment. Throughout the whole experiment you must remain seated and mobile phones and other electronic devices must be switched off. If there are any questions please raise your hand and an experimenter will come to answer your questions in private.

Payment: In this experiment you can earn points. At the end of the experiment you will be paid **£3 as a participation fee** plus **£0.10 for each point earned**. You will be paid in private and in cash**.**

All your decisions are anonymous, so your identity will be kept secret at all times.

Brief description of the task:

At the beginning of the experiment you will be randomly matched with another participant. One of you will have the role of "Buyer" and the other the role of "Seller".

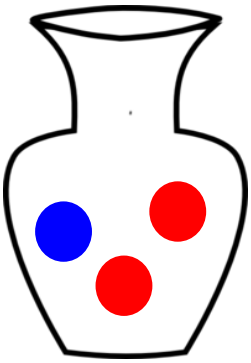
There are two products: a Red product and a Blue product. The Buyer must buy one of these products. The Seller knows the qualities of the two products but the Buyer does not. The final payoff to both participants depends on the buying choice and the quality of the chosen product. A more detailed description of the task is given below.

Detailed description of the task:

The task consists of 2 stages.

Stage 1: *The Seller is informed about the qualities of the two products.*

In this stage, the computer will randomly draw a ball from a computerized urn containing 3 balls: 2 red and 1 blue. The urn is depicted below:



The Seller will observe the colour of the drawn ball, but the Buyer will not.

- If the colour of the drawn ball is **red**, the **Red product** is **better than** the **Blue product.** The Red product is worth 30 points to the Buyer and the Blue product is worth nothing to the Buyer.
- If the colour of the drawn ball is **blue**, the **Blue product** is **better than** the **Red product.** The Blue product is worth 30 points to the Buyer and the Red product is worth nothing to the Buyer.

Stage 2: *The Buyer decides whether to buy the Red product or the Blue product.*

After the Seller observes the qualities of the products, the Buyer will decide between buying either the Red product or the Blue product.

How your earnings are determined:

The Buyer will earn 30 points if they buy the better product. The Seller will earn 30 points if the Buyer buys the Red product. Thus, the point-earnings will be determined as follows.

The Buyer
- earns 30 points if a Red ball is drawn and she buys the Red product;
- earns nothing if a Red ball is drawn and she buys the Blue product;
- earns 30 points if a Blue ball was drawn and she buys the Blue product;
- earns nothing if a Blue ball was drawn and she buys the Red product.

The Seller
- earns 30 points if the Buyer buys the Red product;
- earns nothing if the Buyer buys the Blue product.

This is summarised in the Table below:

	Buy the Red product	**Buy the Blue product**
Red ball drawn	The Seller earns 30 points, The Buyer earns 30 points	The Seller earns 0 points, The Buyer earns 0 points
Blue ball drawn	The Seller earns 30 points, The Buyer earns 0 points	The Seller earns 0 points, The Buyer earns 30 points

Preliminary questions: Before the experiment begins, you will be asked to answer a few questions regarding your understanding of the instructions. The experiment will begin only after all participants have answered these questions correctly.

Final questionnaire: After everyone makes their decisions, you will be asked to fill in a short questionnaire. You will then be paid your earnings in private and in cash. You will receive a £3 participation fee plus £0.10 for each point that you earn.

C.2 SELLER MESSAGES

C.2.1 *CM_13 treatment*

Session	Seller ID	Ball	Message
1	1	blue	The red ball appears better than the blue ball.
1	3	blue	I've got a red ball here. You interested in buying?
1	5	red	begins with the letter R
1	7	red	Both will earn
1	9	red	Buy the red product
1	11	blue	The ball I get is red
1	13	blue	Hey, the red ball was drawn
1	15	red	Buy the red ball - I know it is in my interests for you to choose the red ball but honestly the red ball has been chosen so if you choose red we both get 30 points :)
1	17	red	red ball drawn
1	19	blue	Hello,how are you? I am selling you the product
1	21	red	the red ball is the red ball so pick red and we both get 30 points, It is definitely red i promise, if you buy blue we both get 0 so buy red
1	23	red	It's actually a red ball and it'll be better for both of us if you choose to buy the red ball
1	25	blue	red ball has better quality
1	27	blue	the ball drawn was blue lol so enjoy ur 30
1	29	red	The computer has drawn a red ball (trust me i'll be fair)
2	1	blue	You can buy the red product or the blue product
2	3	blue	Ah lucky got the red ball, So it's a win win right??
2	5	red	Its red. Choose red and we earn together.
2	7	blue	Got a red one

2	9	blue	im not going to lie to you but im also not going to tell you the truth- guess whatever you want and its 50/50 either way
2	11	red	Red ball drawn
2	13	red	Hey, honestly I have drawn the red ball - so please go for the red product and we'll both get 30 points (otherwise we'll both get nothing if you choose the blue one). Cheers
2	15	red	I have drawn a red ball. We both receive a higher pay off if you choose to buy the red ball. If you choose to buy the blue ball then we both get 0 payoff.
2	17	blue	Its all good if you buy the red one we both get 30, these things always go wrong for me so nice I actually get to do the fair option for once ahah. Have a nice day good luck with finals xo, Peace be with you - regardless of your faith
2	19	blue	Buy the Red product, And good luck in life
2	21	blue	blue ball made of plastic and red made of rubber
2	23	blue	The red product is of better quality than the blue product
2	25	blue	We are lucky the ball is a red one, We can both earn 30 points lol I'm actually shaking and nervous about this but just trust me

2	27	blue	Hey, obviously they want to observe how we fool eachother when a blue ball is drawn. But I'm not even gonna try fooling you, I mean in real life, would you fool your customers? No, because they would never come back. At least that's how I see it. Anyway, a blue ball has been drawn, so I recommend you buy the blue product, which I wouldn't recommend, if it wasn't true, because then we both would earn nothing. And no, I'm not gonna be mad at you for earning 30 and me earning 0.
2	29	red	Buy the red ball. Trust me. We just wanna earn the money
3	1	red	red
3	3	red	Red
3	5	red	Buy the red
3	7	blue	Red product
3	9	blue	I have drawn a red ball. I you buy the red ball then we will both win 30 points! If you draw the blue ball we will both earn nothing
3	11	blue	The computer drawn a red ball. Choose the red one so both of us could earn 30 points. or else we would earn nothing
3	13	blue	Which one are you getting?
3	15	blue	I've drawn the red ball, we'll both get 30 points. I'm a person that values honesty more than anything, especially more than 30 points (3 LOL). It's your decision to trust me or not, but I'm being very honest with you here.
3	17	blue	Hi, the Blue product is better than the red product so you can earn points if you select the Blue one while i get no points
3	19	blue	Buy it. If mutually beneficial I will tell you, this one could go either way though
3	21	blue	its a blue ball

3	23	blue	Hiya- I have just drawn a red ball. It makes sense for you to buy the red ball as it will mean we can both get 30 points
3	25	blue	second line of the message
3	27	blue	Hey, the blue ball was picked out
3	29	blue	buy the blue product
4	1	blue	I like leaving this into luck so follow your heart
4	3	blue	the red ball was drawn
4	5	blue	The red ball is considered to be better in all ways and brings happiness all around thus, should be purchased. As a seller, I want to sell the best so it makies everyone happy!
4	7	blue	I think you should buy a red one cause it is better than the blue
4	9	red	pick the red ball and both of us get the benefit.
4	11	blue	the red ball was drawn
4	13	blue	1/3 chance but it came out
4	15	blue	buy the blue ball
4	17	red	it picked the red ball... :) so if you pick too then we can both get the points... either way, you only get the points if you pick the right one, and it really is the red (promise)
4	19	blue	Buy the red one So that we both get points
4	21	red	the red ball is drawn
4	23	blue	you should have a further consideration for buying this, just wait
4	25	red	The RED product is better than the Blue PRODUCT
4	27	blue	The red ball was drawn
4	29	red	Which ball would you like to buy?

C.2.2 *CM_23 treatment*

Session	Seller ID	Ball	Message
1	1	blue	Buy Red
1	3	red	The ball drawn is the red one. With the most outstanding quality! to make sure we are both benefit from this please choose red or otherwise we loose the points given. Cheers.
1	5	red	
1	7	blue	Hi morning. Red ball was drawn. Feel free to decide which ball to buy, as red ball will benefit both of us. Thanks.
1	9	red	Hi, This is a win-win situation, Buy the red one and we both earn 30 points
1	11	red	I think you should buy the red ball, We will both be better off if you do, But we will both be worse off otherwise
1	13	red	red product is better than blue product, the red ball is drawn, you can get points if you choose the right red product
1	15	red	This product is good quality compared to the other in the urn.
1	17	blue	blue ball was drawn so pick the blue product
1	19	red	I have a red ball, For the sake of 3 I see no reason to lie about this to try and fuck you over, if you buy the red then we'll both profit
1	21	red	buy the red ball
1	23	red	the ball is red - 30 points each (not lying, there is a 66% chance the ball is red in the first place) so we both earn 30
1	25	red	It's a red ball, so it's a win-win situation regardless. I will earn 30 points, and you will earn 30 points also. So buy the red product.
1	27	blue	Yo it's red, points all around

1	29	red	Hello. The red ball has been chosen, so if you pick it we will both get 30 points each :)
2	1	blue	Hi! I have drawn a red ball!
2	3	red	Got the Red ball. Have a nice day :)
2	5	red	The ball drawn is red, have to trust me on this! That means both of us get 30 points!! 2/3 chance so I wouldn't be suprised.
2	7	blue	hello! A red ball was drawn. So in order for both of us to earn 30 points, you would have to pick the red ball :)
2	9	blue	It is a red ball
2	11	red	Hi, the red ball has been drawn. I'm going to tell you the truth for the all as it's better for someone to at least get 30 points instead of both of us getting 0. X
2	13	red	the drawn ball is red!
2	15	red	the computer has drawn a red ball. We both earn the points :)
2	17	red	Hi, i am selling the product please slect what u want to buy, red is drawn product
2	19	red	Hey, I drew the red ball - the red product is better than blue!
2	21	blue	computers has drawn a red ball
2	23	red	red product is a better product
2	25	blue	Red is drawn.
2	27	red	I have drawn red
3	1	blue	Hi, The product is blue
3	3	red	it is red, red is good for both of us, please buy!
3	5	red	red is good
3	7	blue	I've got a win-win for both of us here, The red ball has been drawn, 30 points each, Cheeky little extra 3 pound
3	9	blue	The ball drawn is a red ball.
3	11	blue	The chosen ball is blue
3	13	red	buy the red product, red was drawn
3	15	red	a red ball was drawn
3	17	blue	You should buy the Red product

3	19	red	Hi there. so the ball drawn by the computer is the red ball, you have ti trust me here. I am not lying., to trust*, both of us will earn 30 points if u choose the red ball. Thanks
3	21	red	i would like to buy, we both would attain 30 points
3	23	blue	Buy the red product
3	25	blue	i recommend the blue prodduct
3	27	red	the red ball has been drawn out of the urn
3	29	blue	Good Morning!, Regarding to the result of the draw you should choose the red product. Anyway, statistically choosing red randomlz has a higher chance 66%.
4	1	blue	The computer has drawn a red ball.
4	3	red	the ball draw is red, i swear its true *drawn
4	5	red	Hello! The red product is better than the blue product.
4	7	red	ok this is actually great news. Red ball was drawn!!! (2/3 the odds so quite probably) Choose red and we both earn. Please. That's an extra chai latte in starbucks. :)
4	9	red	Hello Buyer, I hope you feel particularly energised and ready to have strawberries
4	11	red	It's red this time gonna be real with you
4	13	red	The aim of the game is to win the most money so therefore I will always be honest with you about what ball is drawn from the urn. However, if it continues to be blue for I would like you to choose the option that gives me the most points. Otherwise we don't trust each other and its a game of luck. The red ball was drawn.

4	15	red	Hey, So, I drew the red ball, I know that you might think that I'm just doing it for my own benefit, But I genuinely drew the red ball, This way we both win, Trust, So either we both win or we both lose!
4	17	red	The red ball was drawn, so go for red and we both get 30 points
4	19	blue	This is a red ball drawn from the urn, if you choose the red we both get 30 points
4	21	red	a red ball has been drawn, so pick the red ball so we both can get 30 points !
4	23	blue	Good afternoon Buyer, Upon analysis of the products available, it was revealed that the red ball is of superior quality compared to the blue ball, therefore I would advise you to buy the red product.
4	25	blue	Hey, buy blue
4	27	red	The red ball is the better product than the blue ball. There will always been a higher probablity of the red ball being chosen as there are 2 red products and one blue product