Three Essays on Beliefs in Behavioral Economics

Dissertation

zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften

(Dr.rer.pol.)

vorgelegt von:

Dominik Bauer

Sektion Politik-Recht-Wirtschaft

Fachbereich Wirtschaftswissenschaften

Konstanz, 2017

Tag der mündlichen Prüfung:	15. Dezember 2017
1. Referent:	Prof. Dr. Urs Fischbacher
2. Referentin:	Prof. Dr. Susanne Goldlücke
3. Referent:	Prof. Dr. Wolfgang Gaissmaier

Danksagung

Ich möchte mich an dieser Stelle bei den Menschen bedanken, die mich in meinen $9\frac{1}{2}$ Jahren in Konstanz und vor allem während der Zeit der Promotion begleitet haben.

Zunächst möchte ich mich bei meinem Doktorvater Urs Fischbacher bedanken. Er hat mir die Promotion überhaupt erst ermöglicht. Er hat für mich und alle anderen exzellente Bedingungen für unsere Arbeit hergestellt, auch wenn ich das –vor allem zu Beginn der Promotion- nicht immer in vollem Ausmaß gemerkt habe. Urs hat uns so gut es ging in finanziellen Dingen und was die Lehre betrifft den Rücken freigehalten. So war es zum Glück kein großer Rückschlag, wenn die ein oder andere Session durch Stromausfälle (oder gar Brände) in den Sand gesetzt wurde. Außerdem konnten wir trotz der großen Gruppe jederzeit einen Termin mit Urs ausmachen um unsere Probleme zu besprechen. Urs, vielen Dank für das alles. Für die Abende im Alti Badi, im Hotel Mohren oder natürlich bei dir zum Fondue. Danke!
Ein großer Dank gilt auch meinen beiden weiteren Betreuern Susanne Goldlücke und Wolfgang Gaissmaier, die diese Arbeit ebenfalls begutachten.

Den Spezialplatz zum Danke sagen bekommt—natürlich—Irenaeus Wolff. Lieber Iren, ich habe auch dir einen großen Anteil meines Erfolges zu verdanken, da du immer mit Hilfe, Anregungen und neuen Treatment-Ideen zur Stelle warst, wenn ich nicht weiter gewusst habe und ich mit dir ausserdem unfassbar oft unsere Projekte besprechen konnte und du stets den Mist den ich produziert habe gelesen, verstanden und verbessert hast (auch wenn ich manchmal mit langen Sätzen und Klammern am Satzende nich einverstanden war). Ich glaube wir haben uns perfekt ergänzt. Vielen Dank Iren!

Der nächste Dank geht an meine Freunde, die glücklicherweise auch meine Kollegen sind. Ohne euch kann ich mir Konstanz, den See, Wochenenden, Skihütten, Fußballabende, Tischtennis, gunf, dune, und das CC nicht mehr vorstellen. Danke an Carl Maier für unzählige Gespräche über Gott und die Welt, für Heimfahrten und grinding-sessions. Danke an Thomas Hattenbach für Kaffee, gunf und Bier. Carl und Thomas haben mich meine gesamte Zeit in Konstanz begleitet und wir haben alle Höhen und Tiefen gemeinsam erlebt. Danke dafür! Danke auch an Fabian Dvorak, Franziska Deutschmann, Nathalie Popovic und Jan Hausfeld als meine Bürokollegen und für Robot Unicorn. Danke an Jan auch für dein Sofa mit Sky (H96 booh!). Danke an Konstantin Hesler, Christian Neumeier, Baiba Renerte, Michal Marencak, Michael Dörsam, David Dohmen, David Grammling, Fadong Chen und Moritz Janas. Danke auch an Kathrin Schmelz, Gerald Eisenkopf und Sebastian Fehrler.
Ich möchte außerdem Alexander K. Wagner danken, der mir erst Spieltheorie beigebracht und mich anschließend doch zum Experimentalökonomen gemacht hat. Bei Ihm habe ich das praktische Experimentieren gelernt und er hat meine Begeisterung für dieses Feld geweckt.
Danke an die Graduate School of Decision Sciences, die mich die ersten drei Jahre finanziert hat. Danke an das Thurgauer Wirtschaftsinstitut, welches die meisten meiner Experimente und mein Abschlussjahr finanziert hat.

Vielen Dank an meine Familie. Danke an Bibi, Siggi, Niki und Patrick. Ihr wart immer für mich da und habt mich mit Rat und Tat unterstützt. Und vor allem: Danke Anja‼!

Contents

Chapter 1

Introduction

1.1 Summary

This dissertation is about beliefs. It is not about religious belief and does not discuss beliefs from the epistemological point of view. Rather, this work is about beliefs how economists use them. When facing a decision, people often do not know the true probabilities of different possible states of the world. It starts from the probability of rain tomorrow, continues with the probability of the DAX reaching 13000 points in a month and ends with the probability of being alive on one's ninetieth birthday. People face uncertainty everywhere. Economic theory assumes, that in such situations people form subjective probabilities and act on those subjective probabilities (*i.e.* beliefs) as if they were the true probabilities. Hence, beliefs are subjective probabilities about unknown states of the world.

In this dissertation, consisting of three papers, I contribute to discussions in the literature about three broad problems, all connected to beliefs. In experimental studies, it has become increasingly popular to transform the originally unobservable beliefs to observables via belief elicitation procedures. Eliciting subjective beliefs can be extremely helpful to test economic models. The first study contributes to the discussion of *how* the belief elicitation procedure should concretely be implemented and identifies biases that are caused by the different methods. The second paper theoretically and empirically studies, why the elicited beliefs are often not consistent with the choices people make. We propose a new from of stochastic choice which is due to uncertainty over beliefs. The third paper contributes to literature on communication in games. By eliciting off-equilibrium beliefs, we are able to show that communication that is theoretically cheap talk, actually conveys information about social preference types.

In the following, I summarize the three studies in more detail.

Chapter 1 - *Biases in Beliefs: Experimental Evidence*

As economists, we often would like to know people's beliefs on which they base their behavior on. Through elicitation procedures, we can make beliefs observable. However, in both psychology and economics, many papers have reported behavioral biases in belief formation that come on top of standard game-theoretic reasoning. In this chapter, we show that the processes involved depend on the way participants reason about their beliefs. When they think about what everybody else or another 'unspecified' individual is doing, they exhibit a consensus bias (believing that others are similar to themselves). In contrast, when they think about what their situation-specific counterpart is doing, they show *ex-post* rationalization, under which the reported belief is fitted to the action and not vice versa. Our findings suggest that there may not be an 'innocent' belief-elicitation method that yields unbiased beliefs. However, if we 'debias' the reported beliefs using our estimates of the different effects, we find no more treatment effect of how we ask for the belief. The 'debiasing' exercise shows that not accounting for the biases will typically bias estimates of game-theoretic thinking upwards.

Chapter 2 - *Belief Uncertainty and Stochastic Choice*

Beliefs are almost always elicited to relate them in some way to behavior. Through the elicited beliefs, economists often want to understand and explain behavior. However, elicited beliefs do not fit to the observed actions more often than we would expect or wish. Belief-action consistency is virtually never at 100% as predicted by standard theory. In this chapter we propose a new form of stochastic choice due to belief uncertainty, which can explain the lack of consistency between beliefs and actions.

We start by the assumption, that people often cannot assign a clear probability to an event but face uncertainty about their subjective probabilities. We model belief uncertainty by assuming that agents' beliefs are characterized by a distribution over subjective-probability distributions that agents cannot access directly. Our model produces stochastic choice because each decision-relevant belief is but one realization out of the distribution over all possible beliefs. Our model predicts that when comparing unknown situations to routine choices, people will make more *ex-ante* suboptimal choices in unknown situations. The model also offers an explanation for experiment participants not playing a best-response to their stated beliefs: participants are uncertain which belief to report or base their decision on, and hence, act on momentaneous 'belief realizations'. In an experiment, we manipulate participants' belief uncertainty. We do so by exogenously manipulating their strategic uncertainty, providing varying levels of information

about historical choice data. In situations in which belief uncertainty is low, observed best response rates are high and increasing in the amount of information we provide. Conversely, high belief uncertainty leads to lower consistency. On top, low belief uncertainty makes participants earn more, even when controlling for the accuracy of their beliefs.

Chapter 3 - *Responder communication in ultimatum bargaining*

In this chapter, elicited beliefs are brought to action. In games of incomplete information and communication, theoretical equilibrium predictions often heavily rely on assumptions about beliefs. Even more, many equilibria rely on fixing certain off-equilibrium beliefs. Therefore, it is often hard to interpret observed actions in such games, because many different preference-belief combinations could potentially have caused the actions. In this study, we bring such a game to the laboratory and also elicit off-equilibrium beliefs to understand observed behavior better. We theoretically and experimentally study responder communication in the ultimatum game. Before the game starts, responders send a structured and non-binding message about the smallest offer, they are willing to accept. Our theoretical results show, that information transmission through communication is limited in equilibrium. In line with this result, we find that responders largely send the message which suggests to split the pie equally, even though their true acceptance threshold is lower in the large majority of cases. Our elicited beliefs show, that i) despite the result before, responders send lower messages than their beliefs would indicate, ii) communication is not cheap talk, because proposers' beliefs are correlated with messages and iii) also proposers deviate from what their beliefs suggest to do. They do not increase offers further for very high messages. The observed discrepancy of beliefs and actions for both player roles is in line with the predictions of inequality aversion, but only if we acknowledge the existence of extreme preference types.

1.2 Zusammenfassung

Diese Dissertation handelt von Beliefs.[1] Sie handelt nicht von religiösem Glauben und beleuchtet Beliefs auch nicht von der epistemologischen Seite. Stattdessen handelt diese Arbeit von Beliefs, wie sie Ökonomen benutzen. Wenn Menschen eine Entscheidung treffen müssen, kennen sie oft nicht die wahren Wahrscheinlichkeiten der unterschiedlichen Umweltzustände. Das beginnt

[1] Ich werde in dieser Zusammenfassung das englische Fremdwort "belief" gebrauchen. Eine gute sinngemäße Übersetzung dafür wäre wohl "Erwartung" oder "subjektive Erwartung". Allerdings sind meiner Meinung nach die englischen Begriffe "expectation" und "belief" nicht zu 100% synonym, deswegen bevorzuge ich das originale Wort "belief".

bei der Wahrscheinlichkeit ob es morgen regnet, geht weiter bei der Wahrscheinlichkeit ob der DAX im nächsten Monat 13000 Punkte erreicht und endet mit der Wahrscheinlichkeit ob man seinen neunzigsten Geburtstag erlebt. Es gibt überall Ungewissheit. Die Ökonomische Theorie nimmt an, dass Menschen in solchen Situationen subjektive Wahrscheinlichkeiten formen und dann so handeln, als wären diese subjektiven Wahrscheinlichkeiten (die Beliefs) die wahren Wahrscheinlichkeiten. Beliefs sind daher subjektive Wahrscheinlichkeiten über unbekannte Umweltzustände. Mit den drei Studien in dieser Dissertation trage ich zu Diskussionen in der Literatur bei, die über drei weitergefasste Probleme handeln, die aber alle mit Beliefs zusammenhängen. Es ist in experimentellen Studien sehr populär geworden, durch Belief-Abfrage-Verfahren die eigentlich nicht beobachtbaren Beliefs in beobachtbare Variablen umzuwandeln. Subjektive Beliefs abzufragen kann sehr hilfreich sein, um Ökonomische Modelle zu testen. Die erste Studie trägt zur Diskussion bei, *wie* das Belief-Abfrage-Verfahren konkret umgesetzt werden soll und identifiziert Verzerrungen, die von den verschiedenen Methoden verursacht werden. Die zweite Studie untersucht - theoretisch und empirisch - warum die abgefragten Beliefs oft nicht konsistent mit dem beobachteten Verhalten der Menschen sind. Wir schlagen eine neue Form von stochastischem Entscheiden vor, welche von Unsicherheit über Beliefs verursacht wird. Und die dritte Studie trägt zur Literatur über Kommunikation in Spielen bei. Wir fragen Beliefs ab, die nicht auf dem Gleichgewichtspfad liegen. Dadurch können wir zeigen, dass Kommunikation, die laut Standardtheorie eigentlich keinen Einfluss haben sollte, tatsächlich Informationen über soziale Präferenzen überträgt.
Im folgenden fasse ich die drei Studien noch etwas ausführlicher zusammen.

Kapitel 1 - *Biases in Beliefs: Experimental Evidence*

Oft würden wir als Ökonomen gerne die Beliefs der Menschen kennen, auf die sie ihre Entscheidungen basieren. Über Belief-Abfrage-Verfahren können wir Beliefs beobachten. Allerdings gibt es in der Psychologie und Ökonomie viele Studien, welche Verzerrungen in der Entstehung von Beliefs gezeigt haben, die zusätzlich zum üblichen spieltheoretischen Denken wirken. In diesem Kapitel zeigen wir, dass die wirkenden Prozesse von der Art abhängen, wie unsere experimentellen Teilnehmer über ihre Beliefs nachdenken. Wenn sie darüber nachdenken, was alle anderen, oder ein nicht näher spezifizierter 'anderer' tut, gibt es einen 'consensus bias' (man glaubt, die Anderen sind einem selbst ähnlich). Im Gegensatz dazu, wenn die Teilnehmer darüber nachdenken, was ihr situationsbezogener Partner tut, tritt ex-post Rationalisierung auf. In diesem Fall werden die berichteten Beliefs der Entscheidung angepasst und nicht anders

herum. Unsere Ergebnisse legen den Schluss nahe, dass es kein 'unschuldiges' Belief-Abfrage-Verfahren gibt, welches unverzerrte Beliefs liefert. Wenn wir unsere geschätzten Effektgrößen der verschiedenen Verzerrungen nutzen um die abgefragten Beliefs zu korrigieren, finden wir keinen Unterschied mehr im Bezug darauf, wie wir die Beliefs abfragen. Diese Bereinigung der Verzerrung zeigt, dass wir das Ausmaß des spieltheoretischen Denkens typischerweise überschätzen, wenn wir die Verzerrungen nicht beachten.

Kapitel 2 - *Belief Uncertainty and Stochastic Choice*

Beliefs werden fast immer dazu benutzt, sie mit dem beobachteten Verhalten zu vergleichen. Oft möchten Ökonomen durch die abgefragten Beliefs das beobachtete Verhalten verstehen und erklären. Allerdings stimmen die Beliefs weniger oft mit dem Verhalten überein, als wir es erwarten oder uns wünschen würden. Die Übereinstimmung zwischen Beliefs und Verhalten erreicht nie die 100%-Marke, die von der Standardtheorie vorhergesagt wird. In diesem Kapitel schlagen wir eine neue Form des stochastischen Entscheidens vor, welches durch die Unsicherheit über Beliefs hervorgerufen wird. Diese neue Form kann die mangelnde Übereinstimmung zwischen Beliefs und Verhalten erklären. Wir nehmen an, dass Menschen oft den verschiedenen Ereignissen keine eindeutigen Wahrscheinlichkeiten zuordnen können. In diesem Fall haben sie Unsicherheit über ihre subjektiven Wahrscheinlichkeiten. Wir modellieren die Unsicherheit über Beliefs indem wir annehmen, dass die Beliefs des Spielers durch eine Verteilung über subjektive Wahrscheinlichkeitsverteilungen charakterisiert werden. Der Spieler hat keinen direkten Zugang zu der Verteilung über Wahrscheinlichkeitsverteilungen. Unser Modell produziert stochastisches Verhalten, weil jeder Entscheidungsrelevante Belief nur eine Ziehung aus der Verteilung aller möglichen Beliefs ist. Unser Modell sagt vorher, dass, wenn man Entscheidungen in unbekannten Situationen mit Routineentscheidungen vergleicht, Menschen in den unbekannten Situationen mehr Fehler machen. Unser Modell liefert auch eine Erklärung, warum das Verhalten von Experimentteilnehmern nicht mit ihren abgefragten Beliefs zusammenpasst: Die Teilnehmer sind sich unsicher, welchen Belief sie angeben sollen, oder auf welchen Belief sie ihre Entscheidung stützen sollen. Daher entscheiden sie auf Basis von spontanen Realisationen ihres Beliefs. In einem Experiment manipulieren wir die Unsicherheit über Beliefs unserer Teilnehmer. Wir variieren die strategische Unsicherheit exogen, indem wir unterschiedlich viel Informationen über Verteilung vergangener Entscheidungen bereitstellen. In Situationen, in denen die Unsicherheit über Beliefs gering ist, stimmen die Beliefs mit dem Verhalten überein. Die Übereinstimmung steigt ausserdem im Umfang der bereitgestellten Informationen. Im Gegen-

satz dazu sinkt die Übereinstimmung, wenn die Unsicherheit über Beliefs groß ist. Unsere Teilnehmer verdienen außerdem mehr, wenn die Unsicherheit über Beliefs klein ist. Dies stimmt auch dann, wenn wir für die Genauigkeit der Beliefs kontrollieren.

Kapitel 3 - *Responder communication in ultimatum bargaining*

In diesem Kapitel werden die abgefragten Beliefs zum Einsatz gebracht. In Spielen mit unvollständiger Information oder mit Kommunikation, beruhen die theoretischen Gleichgewichte oft auf Annahmen über Beliefs. Einige Gleichgewichte beruhen sogar auf Annahmen über Beliefs, die nicht auf dem Gleichgewichtspfad liegen. Dadurch ist es oft schwierig, das beobachtete Verhalten in solchen Spielen zu interpretieren, da viele verschiedene Kombinationen aus Beliefs und Präferenzen zum selben Verhalten führen können. In diesem Kapitel bringen wir ein solches Spiel ins Labor und fragen dabei Beliefs ab, die außerhalb des Gleichgewichtspfades liegen, um das beobachtete Verhalten besser zu verstehen. Wir untersuchen Responder-Kommunikation im Ultimatumspiel in der Theorie und empirisch. Vor dem Spiel können Responder eine strukturierte und nicht bindende Nachricht senden, was das kleinste Angebot ist, das sie gerade noch annehmen würden. Unser theoretisches Ergebnis zeigt, dass die Informationsübermittlung durch die Nachricht im Gleichgewicht eingeschränkt ist. Passend dazu sehen wir im Experiment, dass die meisten Responder die Nachricht senden, welche die gleichmäßige Aufteilung der Auszahlungen vorschlägt. Dies passiert, obwohl die Responder auch weniger als die Hälfte akzeptieren würden. Unsere abgefragten Beliefs zeigen, dass i) trotz des vorherigen Ergebnisses, die Responder niedrigere Nachrichten senden, als es ihre Beliefs vermuten lassen würden, ii) die Kommunikation nicht gehaltlos ist, da die Beliefs der Proposer mit den Nachrichten korrelieren und iii) auch die Proposer von dem Verhalten abweichen, welches man nach ihren Beliefs erwarten würde. Sie erhöhen ihre Angebote nicht weiter für sehr hohe Nachrichten. Die beobachtete Diskrepanz von Beliefs und Verhalten für beide Spielerrollen wird durch die Vorhersagen der Ungleichheitsaversion erklärt. Allerdings nur, wenn man die Existenz von extremen Präferenztypen zulässt.

Chapter 2

Biases in Beliefs: Experimental Evidence[§]

Dominik Bauer Irenaeus Wolff

Graduate School of Decision Sciences & Thurgau Institute of Economics

University of Konstanz, Germany

Abstract: Many papers have reported behavioral biases in belief formation that come on top of standard game-theoretic reasoning. We show that the processes involved depend on the way participants reason about their beliefs. When they think about what everybody else or another 'unspecified' individual is doing, they exhibit a consensus bias (believing that others are similar to themselves). In contrast, when they think about what their situation-specific counterpart is doing, they show *ex-post* rationalization, under which the reported belief is fitted to the action and not vice versa. Our findings suggest that there may not be an 'innocent' belief-elicitation method that yields unbiased beliefs. However, if we 'debias' the reported beliefs using our estimates of the different effects, we find no more treatment effect of how we ask for the belief. The 'debiasing' exercise shows that not accounting for the biases will typically bias estimates of game-theoretic thinking upwards.

JEL classification: C72, C91, D84

Keywords: Belief Elicitation, Belief Formation, Belief-Action Consistency, Framing Effects, Projection, Consensus Effect, Wishful Thinking, Hindsight Bias, *Ex-Post* Rationalization

2.1 Introduction

This paper draws on two lines of research. First, it has become common practice in experimental economics to elicit beliefs along with choices. It is widely acknowledged that elicited

[§]We want to thank Ariel Rubinstein, Yuval Salant, Robin Cubitt, participants at the ESA 2016 European meeting in Bergen, Norway, a seminar audience at the University of Nottingham, the research group at the Thurgau Institute of Economics and the members of the Graduate School of Decision Sciences of the University of Konstanz for their helpful comments. Contact: Chair of Applied Research in Economics, University of Konstanz, Universitätsstraße 10, D-78464 Konstanz, Germany.

beliefs carry useful explanatory power. Accordingly, numerous scholars have worked on the question of how to elicit beliefs from agents that usually are assumed to be expected-utility maximizers.[1] Second, in both psychology and economics, there are countless studies on what biases affect beliefs, probability judgments, and expectations.[2] In this paper, we show that the specific way of asking experimental participants for their beliefs triggers different biases. By prompting our participants to think about their beliefs in different ways and by exposing them to a specific decision environment, we are able to identify the involved psychological processes and hence, the determinants of beliefs. Subjective beliefs play a central role in economic theory. When facing a decision, people often do not know the true probabilities of different states of the world. Standard economic theory assumes that in such situations, people form subjective beliefs and act on those subjective beliefs as if they were the true probabilities (Savage, 1954). Because of this assumption, eliciting subjective beliefs often is extremely helpful to test economic models. The list of examples for this approach is long. Game theorists have tested whether non-equilibrium beliefs can explain non-equilibrium behavior (*e.g.*, Bellemare et al., 2008; Costa-Gomes & Weizsäcker, 2008; Rey-Biel, 2009). Macroeconomists have explained saving and investment decisions by people's beliefs about future income, demand, and inflation (*e.g.*, Guiso & Parigi, 1999; Engelberg et al., 2011). Further examples are development economists studying the adoption of new agricultural technologies (*e.g.*, Delavande *et al.*, 2011b) and health economists studying the reasons for why people engage in activities that put their own health at risk (such as smoking, *e.g.*, Khwaja *et al.*, 2006).[3]

Given belief elicitation has such a broad field of applications, it is crucial that we know how people come up with their beliefs under different circumstances. This is important for interpreting belief-elicitation data from experiments, questionnaires, and surveys, and ultimately for understanding behavior. Therefore, we need to know what biases we bring about by our commonly used elicitation methods. If we trigger specific processes by our elicitation method, we are likely to misinterpret the data systematically.

When studying beliefs, we face a typical conundrum: what theory assumes to be 'the belief' is an unobserved construct in people's heads. If we want to learn anything about it, we have to make beliefs observable. The classical approach in economics to this problem is to observe choices only and recover the unobservables afterwards, for example by invoking the revealed-preference assumption (Samuelson, 1938). However, reconstructing beliefs from choices can

[1] See, *e.g.*, Schotter and Trevino (2014) or Schlag *et al.* (2015).

[2] See *e.g.* Tversky & Kahneman (1974), Gigerenzer & Selten (2001), Gilovich *et al.* (2002), Dhami (2016), Epley & Gilovich (2016), Proto & Sgroi (2017).

[3] For these and further examples, see, *e.g.*, Trautmann & van de Kuilen (2015).

sometimes be very difficult. For example, in numerous games, the same choice can be rationalized by many different beliefs (*e.g.* Manski 2002). For these reasons, an alternative and popular way of making beliefs observable is simply to ask for them in a belief-elicitation procedure.
Now assume we find some systematic bias in the elicited beliefs. What is the origin of the bias? Was the construct in people's heads biased? Or was it the very act of asking, that squeezed 'the belief' into numbers, thereby biasing (only) the report we observe? In our opinion, these questions can be answered only very partially. Trying to tease them apart is beyond the scope of this paper. Hence, when talking about biased beliefs we do not distinguish whether people's true beliefs (which we do not observe) or 'only' the belief reports are biased.

To analyze the interaction between the method we use and the involved psychological processes, we look at three different ways of asking for beliefs which we call 'frames'. The opponent frame asks for beliefs about the participant's matching partner. The random-other frame asks about some other individual who is not the matching partner and the population frame asks for beliefs about the whole population of players. The three ways of asking might trigger different ways of thinking about the belief.
Along with the frames, we consider five different processes as potential determinants of beliefs. *Ex-post* rationalization, wishful thinking, hindsight bias, and consensus bias (projection) are four biases that potentially shape beliefs on top of standard game-theoretic reasoning. In standard game-theoretic models, players first form a belief and then act upon this belief. Under *ex-post* rationalization, this process is reversed: agents first make a choice (by whatever process), and then, they form a belief that justifies their taken action. Wishful thinking has people have more faith in events—including actions taken by others—that would lead to a favorable outcome. Under a hindsight bias, people fail to abstract from their knowledge about the realization of an uncertain event (*e.g.*, their own action, viewed from their opponent's perspective) when evaluating an action that was taken without this knowledge (in this case, their opponent's action). And under a consensus bias, people project onto others what they themselves would do, feel, or think.
The five processes just described will point in the same direction under many circumstances, and in different directions, under others. In this paper, we systematically vary the decision environment as well as the frame of belief elicitation to disentangle the processes and to test which of them play a role under what circumstances.
We run three experiments. In Experiment 1, we show that elicited beliefs display considerable framing differences that also influence observed belief-action consistency. In particular,

we replicate Rubinstein & Salant's (2015) finding that beliefs are closer to participants' own actions under a population frame than under an opponent frame. We conduct an additional treatment, eliciting beliefs under the random-other frame. The results from the random-other treatment shed light on the framing difference between population and opponent frame. The framing difference is caused by the difference 'interaction partner *vs.* another person' and not by whether the question is about 'one person *vs.* several people'.

In Experiment 2, we disentangle the different processes to explain why the observed framing differences occur. We separate consensus bias, hindsight bias, wishful thinking, and game-theoretic reasoning or *ex-post* rationalization. Experiment 2 provides evidence for a consensus effect in the random-other frame. Game-theoretic reasoning and *ex-post* rationalization cannot be distinguished fully, but we find some suggestive evidence for *ex-post* rationalization in the opponent frame.

Experiment 3 corroborates the suggestive evidence from Experiment 2 for *ex-post* rationalization in the opponent frame. All of our experiments provide evidence of game-theoretic reasoning, but none of them provides evidence for a hindsight bias or wishful thinking. Using our estimates on the biases, we can 'recover' participants' hypothetic unbiased beliefs, and provide an estimate for the best-response rate to those 'underlying' beliefs. This exercise suggests that the amount of game-theoretic reasoning typically will be overestimated in many papers in the literature.

In summary, the three experiments provide evidence that different ways of asking for beliefs trigger different specific processes. The population and the random other frame influence beliefs by a consensus bias, and under the opponent frame, participants tend to *ex-post* rationalize their actions via beliefs.

Among other things, this result is important for our understanding of the literature. For example, the consensus bias seems to be closely related to the belief-elicitation method employed by the researchers: it seems that all major studies in economics documenting a consensus bias use the population frame.[4] On the other hand, the opponent frame seems like a natural choice in studies that investigate belief-action consistency and best-response rates.[5] Our systematic investigation shows that the correspondence between population frame and consensus bias is

[4]Selten & Ockenfels (1998), Charness & Grosskopf (2001), Van Der Heijden, Nelissen & Potters (2007), Ellingsen *et al.* (2010, who also use the random-other frame), Engelmann & Strobel (2012), Iriberri & Rey-Biel (2013), Blanco et al. (2014), Danz, Madarász & Wang (2014), Molnár & Heintz (2016), Rubinstein & Salant (2016), Proto & Sgroi (2017).

[5]Costa-Gomes & Weizsäcker (2008), Danz et al. (2012), Hyndman et al. (2012), Hyndman et al. (2013), Manski & Neri (2013), Nyarko & Schotter (2002), Rey-Biel (2009), Sutter et al. (2013), Trautmann & van de Kuilen (2015), Wolff (2015).

more than a mere coincidence.

A consensus bias can be documented only when it stands in contrast to the predictions of standard theory (*i.e.*, when a player wants to choose a different option than her opponent). In such situations, the consensus and *ex-post* rationalization point in opposite directions. However, by the results of this study, a consensus bias is present only under a population or random-other frame. Now consider, as a thought experiment, that the authors documenting the consensus bias had used the opponent frame to elicit beliefs. Not only had they not measured a belief biased towards participants' own actions (because of a consensus bias), they would have measured beliefs biased *away* from participants' actions (because of *ex-post* rationalization under the opponent frame). In other words, had the authors used the opponent frame, they most likely would not have seen a consensus bias but an extraordinarily high proportion of consistent belief-action pairs. This does not mean the consensus bias is not a real phenomenon. However, it may not be as general a phenomenon as the widespread references to it in the literature may suggest.

By the results of our experiment, we recommend to take the substantial framing differences into account in the analysis of existing data or the design of new surveys and experiments. We discuss these consequences in more detail in our concluding section.

2.2 Related Literature

Methods for belief elicitation

The literature has proposed numerous variants for incentives, procedures and mechanisms of belief elicitation (see Schotter & Trevino, 2014, or Schlag *et al.*, 2015, for recent reviews). The large variety of methods and applications also brings about high variation in explanatory power of beliefs for behavior within and across experimental studies.[6] Most of the literature on belief-elicitation methods concentrates on designing different incentive schemes (that is, payoff-rules) and evaluating their performance with respect to belief-action consistency or properness.[7] Additional topics are hedging (Blanco *et al.*, 2010) or the usefulness of second order beliefs (Manski & Neri, 2013). However, systematic investigations of elicitation procedures that are not related to the incentives and their influences on properness and belief-action consistency are rare. There are two noteworthy exceptions.

[6] *E.g.*, in some of the 3×3 games in Costa-Gomez & Weizsäcker (2008), best-response rates are around 51%. On the other end, Manski & Neri (2013) find a best-response rate of 89% in a 2-action Hide&Seek game.

[7] *E.g.*, Armantier & Treich, 2013; Harrison *et al.*, 2014; Hollard *et al.* 2016; Holt & Smith, 2016; Hossain & Okui, 2013; Karni, 2009; Palfrey & Wang, 2009; Trautmann & van de Kuilen, 2015.

Costa-Gomes & Weizsäcker (2008) study belief-action consistency in different generic 3x3 normal-form games. They vary the timing and ordering of belief and action tasks and find no substantial treatment differences. Belief-action consistency is generally low in their study, at around 50%.[8]
In a field study on fishermen in India, Delavande *et al.* (2011a) vary procedural details like the precision with which probabilities can be expressed or how the support of the belief distribution is determined. The authors find that their elicitation results are robust with respect to the methodology.
In the literature, different belief-elicitation treatments usually perform the role of a robustness check. Some papers also pursue a methodologic question, searching for a treatment that truthfully elicits participants' beliefs. Our approach is somewhat different. In this paper, we use different belief-elicitation frames as a treatment to induce different mental representations. These treatments will enable us to learn something about the underlying belief-formation process. Having said that, the results will inevitably speak to methodologic questions as well.

Framing of belief elicitation

Virtually all studies in the literature use the population or the opponent frame, but the specific choice is rarely motivated. As already mentioned, all major studies in economics documenting a consensus effect use the population frame while the opponent frame is the common choice in studies that investigate belief-action consistency and best-response rates.[9] This again underlines the relevance of studying whether observing a consensus bias or particular consistency levels are specific to the belief-elicitation format.
In a completely different context, Critcher and Dunning (2013 and 2014) use the population and the "individual" frame to elicit behavioral forecasts. The individual frame is similar to the opponent and the random other frame in that they ask for the belief about "a randomly selected student… [whose] initials are LB". They find framing differences in judgments of morally relevant behaviors. However, they only elicit beliefs and report a lack of evidence for a consensus effect.

2.3 Determinants of beliefs

We propose that the specific way of asking for beliefs will trigger different processes that shape the belief reports. Hence, the different ways of asking will lead lead to systematic variation

[8]Rubinstein & Salant (2016) also report a "beliefs-first" treatment. Their main effects show up also in this treatment, but less strongly.
[9]See footnotes 3 and 4.

Opponent frame
Object: Single person, the matching partner
"With what probability did your matching partner chose each of the respective boxes of the current set-up?"
Random-other frame
Object: Single person, not the matching partner
"With what probability did a person who is not your matching partner chose each of the respective boxes of the current set-up?"
Population frame
Object: All persons in the session, including the matching partner
"What is the percentage of other participants of today's experiment choosing each of the respective boxes of the current set-up?"

Table 2.1: The three different frames of the belief-elicitation question.

in reported beliefs. In this section, we will first outline the three different ways of asking for beliefs which will also form one of our treatment dimensions in the experiment. Afterwards, we will describe the processes in detail and predict under which of the ways of asking they should be active.

2.3.1 Elicitation frames

The three different questions we use for asking for beliefs are spelled out in Table 2.1. Note that from a standard theory perspective, all three questions are equivalent under a random partner matching protocol.

We will call our different ways of asking for beliefs the elicitation "frames". However, the different ways of asking are more than just frames. It is easy to frame a payoff of 10€ as a gain (*"You receive 10 €"*) or a loss (*"You have 20€, now you loose 10€"*). What we call a "frame" is more than just describing an equivalent outcome in two different ways. Rather, our frames focus on different identities and numbers of people. This pushes participants into thinking about equivalent strategic problems from different perspectives and to focus their thinking on different aspects of the problem.

The opponent frame prompts people to think about their specific interaction partner, although this player is only the one they are randomly matched to out of many other players. In this frame, it seems much more natural to think about the individual incentives of both players and about the strategic interaction they face. In contrast to that, the random-other frame also fo-

	Population	Random Other	Opponent
Game-Theoretic Reasoning	✓	✓	✓
Ex-Post Rationalization (Cognitive Dissonance)	(✓)	-	✓
Wishful Thinking	(✓)	-	✓
Consensus Bias	✓	✓	✓
Hindsight Bias	-	-	✓

Table 2.2: Predictions of which processes are active under which frame.

cuses on an individual person, but since there is no interaction between the players, the strategic aspect is much weaker. The population frame invokes a picture of many other people, most of whom a participant will not interact with. Individual incentives and the strategic aspects may not play as important a role when thinking about the problem on such a "gobal" scale. Rather it seems important what will influence the behavior of the population as a whole.

2.3.2 Processes that shape beliefs

We now describe the different processes and identify under which frame they would be active, if at all. All predictions are summarized in Table 2.2.

Game-Theoretic Reasoning

What beliefs would we expect in the absence of any biases? Beliefs depend crucially on the strategic situation. Put differently, a given game and its payoffs will influence a participant's beliefs and actions. In particular, we would expect beliefs and actions that are consistent with each other, because otherwise the player would be making a mistake in at least one of the two decisions. So, what do we learn when action and belief are consistent? Likely, the agent went through one of two processes: making up a belief and best-responding to it ('game-theoretic behavior'), or first choosing an action by some process and only then making up a belief consistent with the action. This reversed process (action-then-belief) may either be due to the agent's wish to appear consistent (*ex-post* rationalization, Eyster, 2002; Yariv, 2005; Charness & Levin, 2005; Falk & Zimmermann, 2013) or to wishful thinking. We discuss both of these biases in the following.

We expect game-theoretic reasoning to be present under all of the frames, as we are not aware of any study documenting that participants' actions are not positively related to their beliefs.

Ex-Post Rationalization

Agents may want to appear consistent both for external reasons (because they do not want to look like a fool in the eyes of the experimenter) and internal reasons. A prime example of an internal reason is cognitive dissonance (Festinger 1957). In the remainder of this paper, we use "*ex-post* rationalization" as a shorthand for "*ex-post* rationalization due to cognitive dissonance". Under cognitive dissonance, making two inconsistent choices—an action and a belief—would lead to mental unease in the player's mind. In order to avoid such mental unease, the player would adapt her belief to make it fit her taken action.

In light of the above, *ex-post* rationalization should be strongest in the opponent frame: believing that *some other* player chose an option that would be bad for me need not cause cognitive dissonance, because my opponent still might have chosen something else. In contrast, if my *opponent* chose something that would be bad for me given my action, this should indeed cause cognitive dissonance in me. The population frame should be somewhere in between these two cases: the more concentrated my belief in the population frame, the more cognitive dissonance I should experience because the more the population will be representative of my opponent (we do not test this final hypothesis in this paper, though).

Wishful Thinking

A large body of literature studies *unrealistic optimism*, which is described as a tendency to hold overoptimistic beliefs about future events (e.g. Camerer & Lovallo 1999, Larwood & Whittaker 1977, Svenson 1981 or Weinstein 1980, 1989). *Wishful thinking* has been brought forward as a possible cause of unrealistic optimism and has been described as a desirability bias (Babad & Katz 1991, Bar-Hillel & Budescu, 1995). Wishful thinking hence means a subjective overestimation of the probability of favorable events. For example, people believe that things they like are more likely to happen (*cf.* also the closely related idea of *affect* influencing beliefs, Charness & Levin, 2005). Despite the large body of evidence on human optimism (Helweg-Larsen & Shepperd, 2001), there is some doubt about whether a genuine wishful-thinking effect truly exists (Krizan & Windschitl, 2007, Bar-Hillel *et al.* 2008, Harris & Hahn, 2011, Shah *et al.,* 2016). In the context of this study, a player whose belief is influenced by wishful thinking places an unduly high subjective probability on the event that others act such that the player receives a (high) payoff.

We expect wishful thinking to be stronger the more the matching partner is involved, because the desirable outcome depends on this specific person. Hence, wishful thinking should be most

prevalent in the opponent frame, followed by the population frame, and it should be absent in the random-other frame.

Consensus Bias

The *consensus* bias is a prominent phenomenon, intensively studied by psychologists and economists. Tversky & Kahneman (1973, 1974) link it to the *availability heuristic* and the *anchoring-and-adjustment heuristic*. Joachim Krueger gives a very general but engagingly simple description of what the consensus effect means: *"People by and large expect that others are similar to them"* (Krueger, 2007, p. 1). The basic idea of a consensus bias has been studied in many different contexts and it has been given many different names: [false-]consensus effect (Ross, Greene & House, 1977; Mullen *et al.*, 1985; Marks & Miller, 1987; Dawes & Mulford, 1996), perspective taking (Epley *et al.,* 2004), social projection (Krueger, 2007; 2013), type projection (Breitmoser, 2015), evidential reasoning (al-Nowaihi & Dhami, 2015) or self-similarity bias (Rubinstein & Salant, 2016).

For this study, we define the consensus bias broadly, as a psychological mechanism that distorts (reported) beliefs towards a participant's own action. A participant with a belief distorted by a consensus bias reports too high a subjective probability that others choose the same action as herself, relative to the participant's (counterfactual) unbiased belief.

We hypothesize that a consensus bias can occur in all elicitation frames. The expectation, that others are similar to myself, should not depend on whether my matching partner is involved or not. Further, it should not depend on whether thinking about one or many persons.

Hindsight Bias

Under a hindsight bias (Fischhoff, 1975), agents strongly overestimate the probability of an event after the event has materialized. Thus, the hindsight bias is a specific form of information projection (Madarász, 2012). According to information projection, agents cannot abstract from their own information when assessing what other players know. In the special case of the hindsight bias, agents cannot abstract from information that became available only later on when assessing what they or others did before the information became available. Meta-analyses such as Christensen-Szalanski & Willham (1991) and Guilbault et al. (2004) underline the robustness of this effect.

Applied to our setting, a hindsight bias means that players are unable to abstract from the information they have (about their own action) when reporting a belief about others' behavior. Players with a hindsight bias hence form their belief (as if they were) assuming the other players

should have anticipated that the biased player would choose with a very high probability whatever she ended up choosing in actual fact. Therefore, a hindsight bias increases the probability mass placed on the other player(s) playing a best-response to the player's own action.
We expect that a hindsight bias will exclusively occur in the opponent frame, because the hindsight relates to the event that my *matching partner* chooses a best response to my own action. In the random-other frame, the object of belief elicitation does not interact with me. So, this other person will be best-responding to somebody else, which means that the information about my choice should not affect his behavior. Similarly, the population of other players will mostly best-respond to other people, which means the information about my choice will hardly influence their choices.

2.4 Experimental Design

General setup

This paper contains three experiments. We next describe the general setup which all of the experiments have in common as well as the specific purposes of the experiments.
Experiment 1 serves three purposes. First, it replicates the earlier finding that beliefs are closer to participants' own actions under a population frame than under an opponent frame. Second, it showcases the substantial differences the elicitation-frame choice has for interpretations regarding participants' belief-action consistency. And third, it singles out the 'interaction partner *vs.* another person' distinction as the crucial difference between the frames.
Experiments 2 and 3 disentangle the mental processes underlying the findings from Experiment 1. They provide evidence on which of the known biases and processes are important, and when. Experiment 2 separates the consensus bias, hindsight bias, and wishful thinking from game-theoretic reasoning and *ex-post* rationalization. Experiment 3 separates ('*ex-ante*') game-theoretic reasoning from *ex-post* rationalization.

In particular for Experiments 1 and 2, it is crucial to control participants' preferences because we want to interpret belief-action consistency. Abstracting from stochastic choice and stochastic belief reports (see, *e.g.*, Bauer & Wolff, 2017), belief-action inconsistencies can happen for two reasons: (i) the researcher may have mis-specified the participants' utility function, and (ii) participants may have a bias in their belief reporting. This paper focuses on participants' biases in the reporting of beliefs. In contrast to some of the earlier literature, we choose games in which

the predictions do not change for any of the well-documented deviations from risk-neutral payoff maximization. We thereby rule out mis-specification of participants' utility function as a reason for belief-action inconsistency. In particular, non-neutral risk and loss attitudes and social preferences do not matter for the predictions in the games we chose.[10]

In Experiments 1 and 3, participants face a series of 24 one-shot, two-player, four-action pure discoordination games. Players get a prize of 7€ if they choose different actions and nothing, otherwise. Participants play the discoordination games on different sets of labels such as "1-2-3-4", "1-x-3-4", or "a-a-a-B", with randomly changing partners, and without any feedback in between.[11] In Experiment 2, we use the same general setup. However, participants play one-shot "to-your-left-games" (Wolff, 2017), in which a player gets a prize of 12€ if he chooses the action immediately to the left of his opponent. The game works in a circular fashion, so that choosing "4" against a choice of "1" by your opponent would make you win the 12€ in a "1-2-3-4" setting.[12]

Along with every choice in the game, we elicit probabilistic beliefs in every period, incentivizing the belief reports via a Binarized-Scoring Rule (McKelvey & Page, 1990, Hossain & Okui, 2013). In the belief-elicitation task, subjects could earn another 7€. The Binarized-Scoring Rule uses a quadratic scoring rule to assign participants lottery tickets for a given prize. The lottery procedure accounts for deviations from risk neutrality and, under a weak monotonicity condition, even for deviations from expected utility maximization (Hossain & Okui, 2013). Hence, we control for participants' (risk) preferences also in the belief task.

The exact framing of the belief-elicitation question is subject to treatment variation as described in Section 2.3.1. At the end of the experiment, we randomly select two periods for payment. In one period, we pay the outcome of the game and in the other period, the belief task.

Procedures

The experiments were programmed using z-Tree (Fischbacher, 2007) and conducted in the Lake-Lab at the University of Konstanz. In total, we recruited 301 participants using ORSEE (Greiner, 2015). All sessions lasted between 60 and 90 minutes. See Appendix 2.B for the instructions.

[10] More precisely, social preferences do not matter in Experiments 1 and 3 unless participants have so spiteful preferences that they prefer both participants receiving nothing to both receiving the same positive amount of money. This case should happen so rarely that we abstract from it. In Experiment 2, social preferences do not matter as long as people are not ready to burn own money for the sake of equality (a condition that already Fehr & Schmidt, 1999, impose).

[11] For the full list of label sets, see Table A1 in the appendix. All participants went through the same order of sets. We chose the varying sets to keep up participants' attention.

[12] The difference in payoffs is meant to reduce expected-earnings differences accross experiments. In a discoordination game, (both) participants are likely to win fairly often, while in the "to-your-left-game", participants will win at a much lower rate.

2.5 Experiment 1: Framing effects and belief-action consistency

Rubinstein & Salant (2015) find in a chicken-game experiment that beliefs are closer to participants' own actions under a population frame than under an opponent frame. In Experiment 1, we replicate the effect for a pure discoordination game. Note that changing the population frame to an opponent frame changes three things at a time. The first change is that the opponent frame asks for our belief about the person we are currently interacting with, while the population frame is mostly or even fully about "irrelevant" others (interaction partner *vs.* another person). The second change is that the opponent frame is about one person, while the population frame is about several people. Hence, the target is a different statistical object. And finally, because the targets are different, the absolute level of incentives is different.[13]

Following Rubinstein & Salant's (2015) argument and our own intuition, we conjectured that the relevant difference was the difference "interaction partner *vs.* another person". To test this conjecture, we included a third belief-elicitation frame, the random-other frame. This frame asks about the choice probabilities of a randomly drawn 'non-interaction-partner'. This is a *ceteris-paribus* comparison, as both the level of incentives and the number of observations in the target remain unchanged. We analyze the data of 145 participants from Experiment 1.[14] We elicited beliefs directly after each action.

Results of Experiment 1

Figure 2.1 summarizes beliefs and belief-action consistency for the three frames in the discoordination game. For the analysis, we aggregate the data on the individual level across all periods. For each participant, we look at the probability mass in the reported belief on the participant's own action in the corresponding game, averaged across all 24 periods. This is the average subjective probability that the participant did *not* discoordinate. This procedure yields one independent observation per participant. Similarly, we compute the best- and 'worst-response' rate to beliefs for each participant individually. A worst-response means that the participant

[13] To see this, think about the case that a participant knows the distribution of others' choices exactly. Then by design, it is optimal for the participant to report the true probabilities under either frame. However, this means that she will obtain the 'belief prize' with certainty under the population frame because the reported distribution is compared to the true distribution. Under the opponent frame, she will obtain the prize with a much lower probability, because her report is compared to *one realization* instead of being compared to the full distribution. In our design, the probability of receiving the prize when (optimally!) reporting the true choice distribution can be as low as 62.5% (under a uniform choice distribution).

[14] We exclude one participant from Experiment 1 who always reported a 100% belief of not having discoordinated. This participant probably tried to hedge, but did not understand that hedging was impossible.

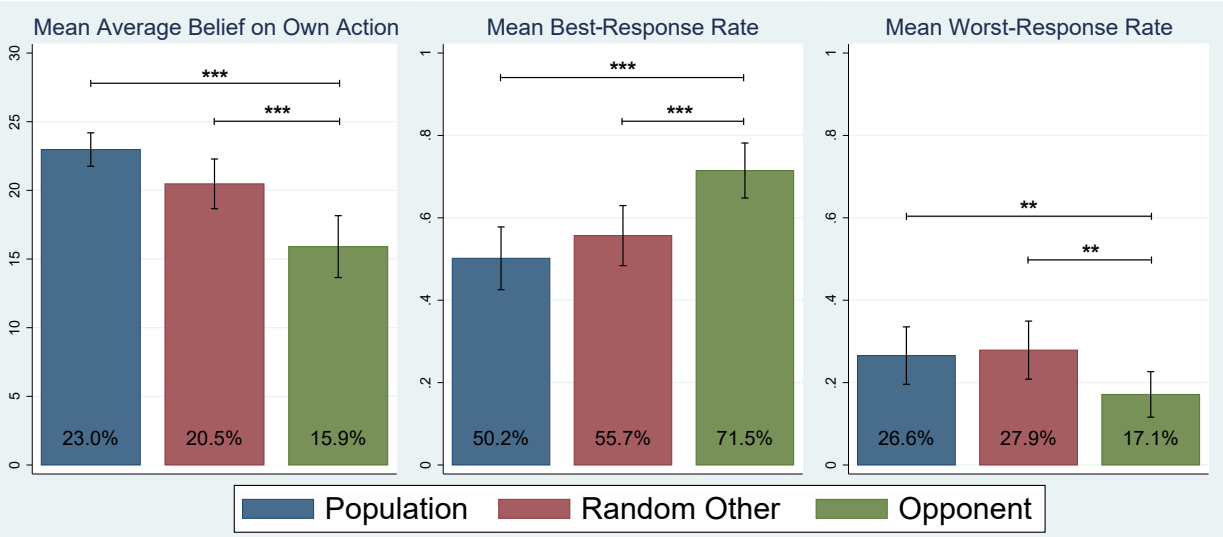

Figure 2.1: Beliefs and consistency in Experiment 1. Error bars indicate 95% confidence intervals. Rank-Sum tests: *** $p < 0.01$, ** $p < 0.05$. For all tests, the data is aggregated on the individual level across all periods, yielding one independent observation per participant.

chooses the action his opponent is most likely to choose, as judged by the participant's reported belief.

The mean average belief on the participant's own action (Figure 2.1, left panel) is significantly higher in the population frame and the random-other frame compared to the opponent frame (rank-sum tests, population/opponent: $p < 0.001$ and random-other/opponent: $p < 0.001$). The effect is strong enough to impede consistency: compared to the opponent frame, the average observed best-response rate is lower (mid panel, $p < 0.001$ and $p = 0.004$) and the average worst-response rate is higher (right panel, $p = 0.026$ and $p = 0.019$) in the population frame and the random-other frame.[15] The reduction in the best-response rate of more than 20 percentage-points and a 9.5 percentage-point increase in the worst-response rate in the population frame are considerable effect sizes. Note that the worst-response rate differs by more than 50% of the rate in the opponent frame.

For a more detailed picture of the results, we depict cumulative distribution functions of the same data in Figure 2.2. Own-action probabilities in the population frame second-order stochastically dominate those in the opponent frame and the distributions differ significantly according to a Kolmogorov-Smirnov test ($p < 0.001$). This effect again carries over to consistency: The best-response rate distribution in the opponent frame first-order stochastically dominates the respective distribution of the population frame and the distributions differ significantly ($p = 0.001$). Similar results hold when comparing the distributions of the opponent and the

[15] There is no significant difference between population and random-other frame. Rank-sum tests, beliefs: $p = 0.146$, best-response rates: $p = 0.237$, worst-response rates: $p = 0.822$.

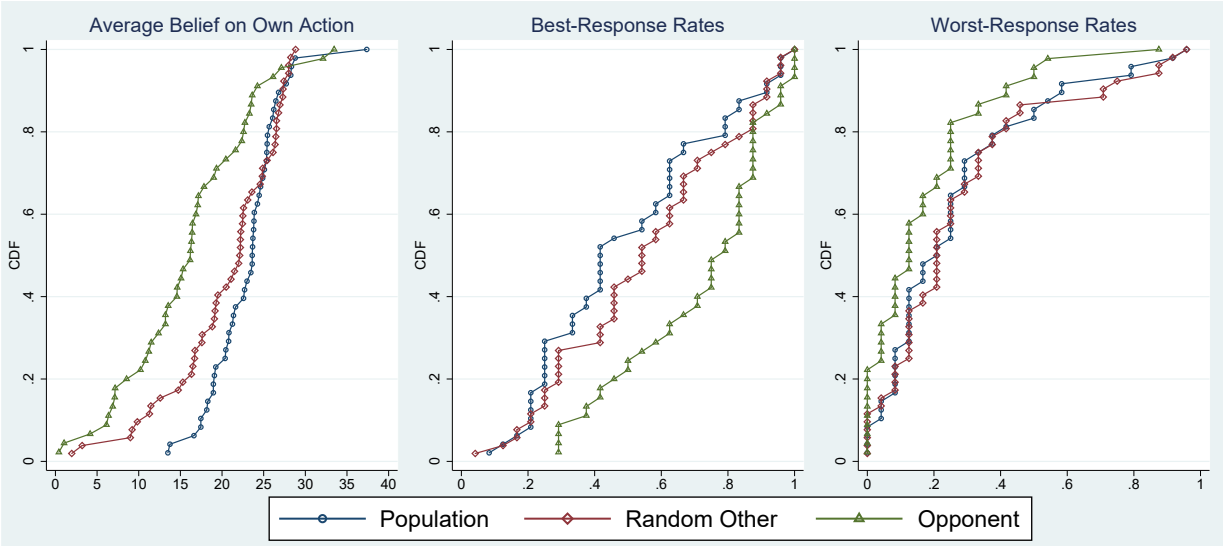

Figure 2.2: Cumulative distributions of individual belief and consistency data in Experiment 1 across frames.

random-other frame (beliefs: $p = 0.002$, best-response rates: $p = 0.008$).[16]

Summary of Experiment 1

The results of Experiment 1 show a considerable framing effect in belief elicitation. Although the questions in all frames are theoretically equivalent (up to the absolute level of incentives), reported beliefs differ substantially across frames. Most notably, beliefs differ in the *ceteris-paribus* comparison between the opponent and the random-other frame, where we vary only the identity of the target participant. Additionally, the differences in reported beliefs influence observed best- and worst-response rates and hence affect the interpretation of actions and beliefs by the experimenter. What Experiment 1 does not show is whether the differences between the frames occur because there is (more) consensus under the population and random-other frames, or because there is (more) hindsight bias, wishful thinking, game-theoretic reasoning, or *ex-post* rationalization under the opponent frame. To disentangle these processes, we need Experiments 2 and 3.

2.6 Experiment 2: Disentangling the processes

Experiment 2 is designed to explain why the framing effects documented in Experiment 1 occur and to disentangle the influences of a consensus bias, hindsight bias, wishful thinking, and

[16]The distributions of the population and random-other frames do not differ significantly. Kolmogorov-Smirnov tests, beliefs: $p = 0.174$, best-response rates: $p = 0.305$. There is also no significant difference between the distributions of worst-response rates across frames (all $p > 0.122$).

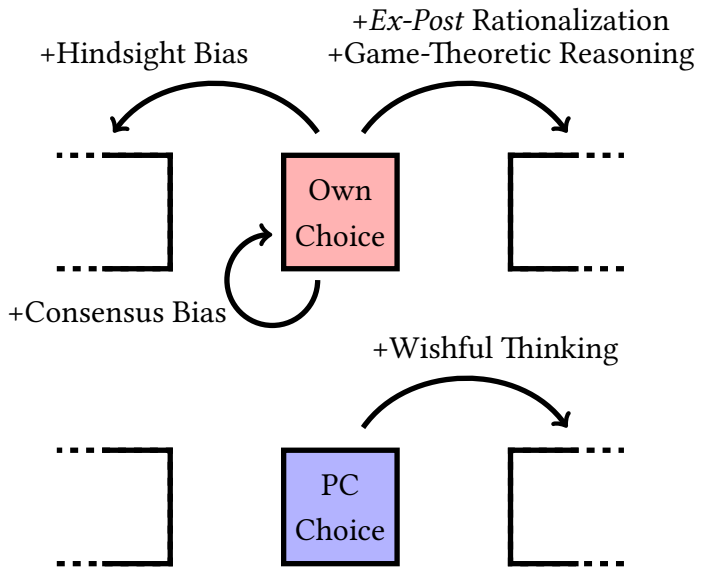

Figure 2.3: Predictions of the candidate processes in the to-your-left game with implementation errors in the case of an implementation error. We color example choices and indicate by arrows the predictions of the four candidate effects that depend on the relative position of the choices.

game-theoretic reasoning/*ex-post* rationalization, which is not possible in the standard discoordination game. For this purpose, we use the "to-your-left game", in which a player wins a prize of 12€ if she chooses the option to the immediate left of the other player's choice (with the right-most option winning against the left-most option).

Predictions for Experiment 2

Figure 2.3 visualizes the predictions of our candidate processes in experiment 2. Because the game is circular, only the relative position of the respective box matters and not the actual position.

In the to-your-left game, a consensus bias still would increase the belief-probability mass participants place on their own actions. A hindsight bias would increase the probability mass on the option immediately to the left of participants' choices, because in hindsight, it would be obvious what the participant's opponent should have chosen in response to the participant's own action. Game-theoretic reasoning, *ex-post* rationalization, and wishful thinking, on the other hand, would increase the probability mass on the option immediately to the right of participants' chosen actions. To separate wishful thinking from game-theoretic reasoning and *ex-post* rationalization, we introduce random implementation errors. In every period, after the participant chooses one of the boxes, there is a 50% probability that the computer changes the participant's decision. If the computer alters the decision, the computer chooses each box with

equal probability (including the participant's chosen box). If the computer changes the decision, the computer's choice is used to determine the game payoff of the participant and of her interaction partner. However, the belief elicitation following each action always targets the other participants' original choices, not the implemented ones. This means that when the computer changes the decision, wishful thinking would increase the probability mass of the option to the right of the implemented decision. In contrast, game-theoretic reasoning and *ex-post* rationalization still mean a higher probability mass on the option to the right of the participant's originally chosen option.[17]

We elicit beliefs directly after each action and 70 participants took part in Experiment 2. We use only the random-other and opponent frames since they provide the most conservative treatment comparison by changing only the identity of the target.

Results of Experiment 2

We analyze the data from Experiment 2 with a dummy regression reported in Table 2.3. The dependent variable is the reported belief on a single box. Every participant reports 24 Periods $\times$ 4 Boxes = 96 beliefs on single boxes. We regress the beliefs on a set of dummies, indicating whether the particular belief can be influenced by a consensus bias, wishful thinking (WT), hindsight bias, or game-theoretic reasoning/*ex-post* rationalization (GT/EPR) according to the predictions above. Further, we use a frame dummy which is equal to 1 in the random-other frame and 0 in the opponent frame. The constant of this regression is a neutral belief where all dummies are zero. Hence such a belief is unaffected by our candidate effects. Model 1 uses all observations where the participant made the ultimate decision.[18] Wishful thinking and GT/EPR cannot be distinguished for the undistorted choices, as both load on the probability to the immediate right of the participant's choice. We hence have to use two separate regressions for the situations with and without implementation error because by design, the interaction GT/EPR $\times$ WT is perfectly collinear with the implementation error.

Model 1 shows that there is a consensus bias, but only in the random-other frame. There is no evidence for a hindsight bias. Further, probabilities to the right of the chosen option (influenced by GT/EPR and/or WT) are twice the size of a neutral belief. This huge effect in the opponent

[17]Note that in some cases, depending on which box the computer selected, two different processes would increase the belief-probability mass on the same option. We control for this in the analysis.

[18]All results in Model 1 are robust to adding trials to the sample in which the computer decided but happened to choose the same action as the participant. A regression that controls for trials in which the computer randomly implemented the same option as the participant detects no significant differences between the two situations. The regression has an additional dummy for 'same choice by computer' which we interact with all six exogenous variables from Model 1. We report the regression in Table A1 in the Appendix.

Single Belief	Model 1	Model 2
False consensus	-0.127	0.701
	(2.132)	(1.980)
False consensus × Random-Other Frame	7.677***	-0.043
	(2.802)	(2.165)
Hindsight Bias	-1.729	-1.211
	(1.819)	(1.799)
Hindsight Bias × Random-Other Frame	1.481	-1.839
	(2.070)	(2.195)
Belief to the right (GT/EPR & WT)	19.353***	
	(3.436)	
Belief to the right (GT/EPR & WT) × Random-Other Frame	-6.650*	
	(3.924)	
GT/EPR		8.690***
		(2.529)
GT/EPR × Random-Other Frame		-2.257
		(2.588)
Wishful thinking (WT)		-0.451
		(1.081)
Wishful thinking (WT) × Random-Other Frame		2.364
		(2.542)
Neutral Belief (constant)	20.301***	23.282***
	(1.031)	(0.870)
Implementation error	No	Yes
Number of Observations	3332	2532
Number of Clusters	70	70
R^2	0.1254	0.0389

Table 2.3: Linear dummy regressions of single belief elements. Standard errors in parentheses clustered on subject level. Asterisks: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

frame is reduced when using the random-other frame. We want to argue that this reduction is indirect evidence of *ex-post* rationalization.

Ex-post rationalization should occur exclusively (or at least to a much larger degree) in the opponent frame: believing that *some other* player chose an option that would be bad for us need not cause cognitive dissonance, because our opponent still might have chosen something else. In contrast, if we state a belief that our *opponent* chose something that would be bad for us given our action, this should indeed cause cognitive dissonance in us. Therefore, the coefficient of "Belief to the right" (with Frame = 0) should capture the added effects of game-theoretic reasoning and *ex-post* rationalization. The "Belief to the right" in the random-other frame (Frame = 1) should capture mostly game-theoretic reasoning only and no (or less) *ex-post* rationalization. Hence, the interaction effect "Belief to the right × Frame" provides an estimate for the differential effect of *ex-post* rationalization.

Like in Experiment 1, the average best-response rate is higher in the opponent frame than in the random-other frame when the computer does not change the decision (opponent: 62.1%, random other: 45.2%, rank-sum test $p = 0.006$).[19]

Model 2 includes all decisions where the computer really changed the participant's decision.

[19]The difference in worst-response rates is not significant. Opponent: 20.9%, random other: 22.8%, $p = 0.780$

Hence, Model 2 includes all observations in which the computer decided and did not choose the same action as the participant. There is no more consensus effect in either frame. Also, there is no evidence for wishful thinking or a hindsight bias. However, GT/EPR loads on beliefs to the right of the participant's decision also in the randomly altered trials. Further, (neutral) beliefs are closer to uniformity in the random-action trials. The results of Model 2 are robust to including all possible remaining dummy interactions.[20]

Estimates of unbiased beliefs

The results in Table 2.3 also give evidence on the size of the respective biases. Having quantified the biases, we are able to reconstruct an estimate for participants 'unbiased' beliefs. We do this correction for all observations used in Model 1. To do so, we subtract the estimated coefficients for the biases from participants' reported beliefs whenever indicated by the respective dummy variables. Subsequently, we re-scale the beliefs to 100%. This procedure yields estimates for unbiased beliefs only on the average level because participants might differ, for example, in how strongly they project their own behavior onto others. Further, we exclude beliefs with multiple best responses and extreme beliefs (that place 100% probability mass on one box) from the correction. Uniform or extreme beliefs are likely to be formed by some alternative process, where the biases do not apply.[21] It hence does not make sense to correct for the biases in these cases. For consistency, we re-run Model 1 in Table 2.3 excluding beliefs with multiple best-responses and extreme beliefs for the correction. The estimation results are similar to the results in Table 2.3 and reported in Table A2 in the Appendix.[22]

We correct beliefs for the consensus effect and hindsight bias, and depending on the frame. As already mentioned, we interpret the coefficient of (Belief to the right $\times$ Frame) as the effect size of *ex-post* rationalization in the opponent frame. We hence correct beliefs for this coefficient as well, but *not* for our estimate of Game Theoretic thinking (which would be 'Belief to the right' + 'Belief to the right $\times$ Frame'). We then compare actual decisions and corrected beliefs, and compute the best-response rate under the hypothetic 'unbiased' beliefs. We do this for every participant separately.

As we have shown above, the original best-response rates differ across frames.[23] However, the

[20]The interactions are: (False consensus $\times$ Wishful thinking), (False consensus $\times$ Wishful thinking $\times$ Frame), (Hindsight Bias $\times$ Wishful Thinking) and (Hindsight Bias $\times$ Wishful Thinking $\times$ Frame).

[21]For example, it seems very unlikely that people hold unbiased beliefs very often that are exactly uniform *after* the biases play out.

[22]The following results continue to hold when we use the unrestricted estimates in Table 2.3 to correct beliefs.

[23]The original best-response rates differ also when using only observations with unique best-responses and which are not extreme (opponent: 55.1%, random other: 42.3%, rank-sum test $p = 0.071$).

corrected average best-response rates do no longer differ significantly across frames (opponent: 46.2%, random other: 44.8%, rank-sum test $p = 0.959$). This result suggests that we can 'debias' the reported beliefs to estimate the true amount of game-theoretic thinking in the to-your-left game. In this perspective, the original best-response rates are biased upwards in the opponent frame (signed-rank test, $p < 0.001$) and biased downwards in the random-other frame ($p = 0.013$).

Discussion of Experiment 2

We interpret the results in the following way: there is a consensus bias in the random-other frame. There is 'game-theoretic reasoning' in both frames, but it is stronger in the opponent frame. We argue that this difference is due to *ex-post* rationalization, which is less important or absent in the random-other frame. Finally, a hindsight bias does not seem to play a role.

As in Experiment 1, the framing differences in Model 1 affect measured belief-action consistency, with higher observed best-response rates under the opponent frame compared to the random-other frame. Using the estimates, we can correct for the observed biases to find participants' hypothetic 'true beliefs before reporting' and assess the amount of game-theoretic reasoning in the game. Our results suggest that this is indeed possible. The framing difference vanishes under the corrected beliefs, which suggests that we did not miss any process that would affect beliefs differentially in the two treatments. The estimated 'true' best-response rates of about 45% suggest the degree of game-theoretic reasoning may be over-estimated in many of the existing studies.

When the computer overrides participants' decisions, only a certain degree of game-theoretic reasoning survives in the reported beliefs: also in such cases, participants *on average* seem to report beliefs that make sense given their actions, despite the fact that beliefs are closer to uniformity.[24] However, there are no more significant framing differences in beliefs or best-response rates with implementation errors. It seems as if the random implementation error detaches participants to a large degree from the action choice altogether. We also do not see any evidence for wishful thinking, even though wishful thinking does not relate to the chosen action.

Experiment 2 was able to disentangle consensus bias, hindsight bias, and—albeit with a caveat—wishful thinking from game-theoretic reasoning/*ex-post* rationalization. The results hint towards overestimated observed best-response rates under the opponent frame, mainly due to

[24]The reduced average difference to uniformity is only very partially due to a difference in the prevalence of uniform beliefs: under implementation errors, 5% of the reported beliefs are uniform, and without errors, 4%.

ex-post rationalization, and underestimated best-response rates in the random-other frame due to a consensus effect. However, the evidence with respect to the discrimination between game-theoretic reasoning and *ex-post* rationalization is only suggestive. To disentangle these two aspects, we need Experiment 3.

2.7 Experiment 3: Identifying *ex-post* rationalization

In Experiment 3, we eliminate the potential for *ex-post* rationalization in the opponent frame by asking participants about their beliefs (directly) *before* they make their choice in the discoordination games from Experiment 1.[25] Comparing the own-action probabilities from this treatment to the corresponding probabilities from Experiment 1 yields an estimate for the importance of *ex-post* rationalization. We can interpret the probability difference in this way because we already know from Experiment 2 that both the consensus effect and wishful thinking do not seem to play a role under the opponent frame. As an additional benchmark, we also ran two sessions under the random-other frame. Under this frame, we expect there to be no difference between Experiment 1 and Experiment 3 (as stated above, we see little scope for *ex-post* rationalization in the random-other frame). 86 subjects participated in Experiment 3.

Results of Experiment 3

The results in Figure 2.4 show that removing the potential for *ex-post* rationalization indeed changes the own-action probabilities in participants' reported beliefs: under the opponent frame—the frame under which we would expect *ex-post* rationalization—average own-action probabilities are roughly four percentage points (or 25%) higher when beliefs are elicited before actions compared to when they are elicited after the action (rank-sum test, $p = 0.028$). In contrast, under the random-other frame (where we argued *ex-post* rationalization should play no role) there is no difference ($p = 0.742$), which is in line with the results of Rubinstein & Salant (2016). We interpret the results as additional evidence for *ex-post* rationalization in the opponent frame.

[25] Giving a belief and then choosing an action that fits this belief seems rather unintuitive: we may well choose an action without forming a belief in the standard setup, but once we form a belief (as in the first stages of Experiment 3), there does not seem to be a good reason to form yet a different belief that we then contradict out of a taste for consistency.

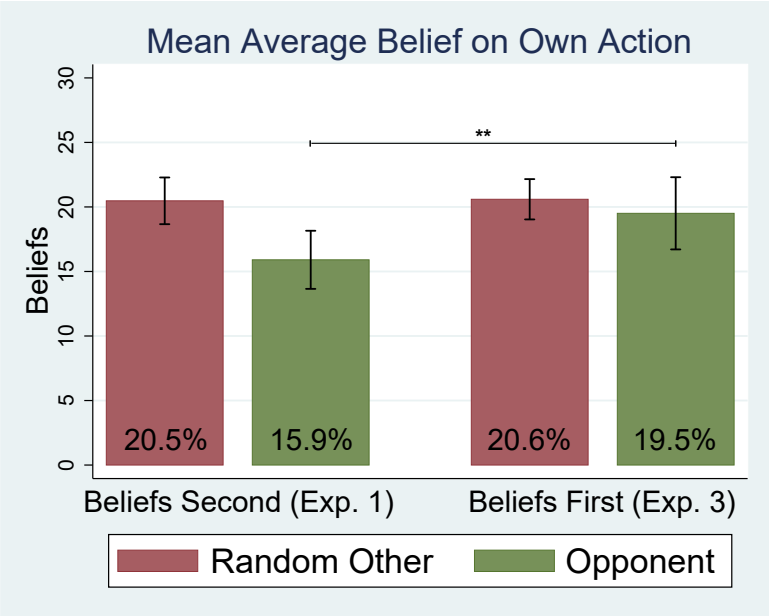

Figure 2.4: Beliefs in the Beliefs-First and the Beliefs-Second treatments. Error bars indicate 95% confidence intervals. Rank-Sum tests: ** $p < 0.05$. For all tests, the data is aggregated on the individual level across all periods yielding one independent observation per participant.

2.8 Conclusion

This paper uses several experimental manipulations to study under which circumstances game-theoretic thinking, *ex-post* rationalization, hindsight bias, wishful thinking, and a consensus bias influence a person's reported belief. Eliciting beliefs in a question targeting people who are not the participant's current interaction partners causes beliefs to be influenced by a consensus bias. A participant with such a belief reports a high subjective probability that others choose the same action as herself. When the question focuses on the participant's current matching partner, there is evidence of *ex-post* rationalization. Under *ex-post* rationalization, the reported belief is fitted to the action and not vice versa. There is no evidence of a hindsight bias or wishful thinking but substantial game-theoretic thinking in all conditions. This means that reported beliefs are consistent with behavior on average. However, the systematic variation in beliefs affects belief-action consistency in predictable ways.

The findings suggest that there may not be an 'innocent' belief-elicitation method. In this study, participants faced a comparatively strong monetary incentive to report their true beliefs. Moreover, we incentivized the belief reports by a state-of-the-art mechanism that is proper even for people who do not comply with expected-utility maximization. And still, we do not seem to be able to find a way of asking for a belief that leads to an unbiased belief-report.

However, by correcting beliefs for the biases, we are able provide an estimate for participants' unbiased beliefs. Using the 'debiased' beliefs, we calculate a 'debiased' best-response rate. The 'debiased' best-response rates suggest that we included all relevant biases and processes as, after the correction, there is no framing difference left to explain. The 'debiased' best-response rate also provides a strong indication that many of the papers in the literature may have over-estimated the degree of game-theoretic reasoning present in economic experiments.

On a methodologic note, our findings are important for experimental researchers who wish to elicit beliefs. The choice of method brings about systematic differences in results. For example, our findings are able to shed some light on why studies documenting a consensus bias all seem to use a population frame, while studies that are after consistency use the opponent frame. Moreover, our findings can inform also other applied researchers: in surveys about inflation, future demand, and other important indicators, reported expectations are likely to be biased. First, a manager might *ex-post* rationalize a recent investment decision by reporting favorable expectations. Hence, researchers will have to control for major question-related recent investment decisions (and be it the decision *not* to invest). On the other hand, when asked for the outlook of a typical company of the same branch, the manager might project an unfavorable situation of the manager's own company onto other enterprises, downplaying the importance of other relevant indicators. These considerations provide support for the necessity of taking into account the effects of belief biases in any survey, questionnaire, or experiment that asks people for their beliefs.

References

al-Nowaihi, A., & Dhami, S. (2015). Evidential Equilibria: Heuristics and Biases in Static Games of Complete Information. *Games*, 6(4), 637-676.

Armantier, O., & Treich, N. (2013). Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. *European Economic Review*, 62, 17-40.

Babad, E., & Katz, Y. (1991). Wishful thinking—against all odds. *Journal of Applied Social Psychology*, 21(23), 1921-1938.

Bauer, D., & Wolff, I. (2017). Belief uncertainty and stochastic choice. *Mimeo*.

Bar-Hillel, M., & Budescu, D. V. (1995). The elusive wishful thinking effect. Thinking & Reasoning, 1(1), 71-103.

Bar-Hillel, M., Budescu, D. V., & Amar, M. (2008). Predicting World Cup results: Do goals seem more likely when they pay off? *Psychonomic Bulletin & Review*, 15(2), 278-283.

Bellemare, C., Kröger, S., & Van Soest, A. (2008). Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities. *Econometrica*, 76(4), 815-839.

Blanco, M., Engelmann, D., Koch, A. K., & Normann, H. T. (2010). Belief elicitation in experiments: is there a hedging problem?. *Experimental Economics*, 13(4), 412-438.

Blanco, M., Engelmann, D., Koch, A. K., & Normann, H. T. (2014). Preferences and beliefs in a sequential social dilemma: a within-subjects analysis. *Games and Economic Behavior*, 87, 122-135.

Breitmoser, Y. (2015). Knowing me, imagining you: Projection and overbidding in auctions. *Working paper*, accessed 2017/09/06, https://mpra.ub.uni-muenchen.de/62052/

Camerer, C., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *The American Economic Review*, 89(1), 306-318.

Charness, G., & Grosskopf, B. (2001). Relative payoffs and happiness: an experimental study. *Journal of Economic Behavior & Organization*, 45(3), 301-328.

Charness, G., & Levin, D. (2005). When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect. *The American Economic Review*, 95(4), 1300-1309.

Christensen-Szalanski, J. J., & Willham, C. F. (1991). The hindsight bias: A meta-analysis. *Organizational Behavior and Human Decision Processes*, 48(1), 147-168.

Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies*, 75(3), 729-762.

Critcher, C. R., & Dunning, D. (2013). Predicting persons' versus a person's goodness: Behavioral forecasts diverge for individuals versus populations. *Journal of Personality and Social Psychology*, 104(1), 28.

Critcher, C. R., & Dunning, D. (2014). Thinking about Others versus Another: Three Reasons Judgments about Collectives and Individuals Differ. *Social and Personality Psychology Compass*, 8(12), 687-698.

Danz, D. N., Fehr, D., & Kübler, D. (2012). Information and beliefs in a repeated normal-form game. *Experimental Economics*, 15(4), 622-640.

Danz, D. N., Madarász, K., & Wang, S. W. (2014). The Biases of Others: Anticipating Informational Projection in an Agency Setting. *Working Paper*, accessed 2017/06/06, http://works.bepress.com/kristof_madarasz/42/

Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment?. *Organizational Behavior and Human Decision Processes*, 65(3), 201-211.

Delavande, A., Giné, X., & McKenzie, D. (2011a). Eliciting probabilistic expectations with visual aids in developing countries: how sensitive are answers to variations in elicitation design? *Journal of Applied Econometrics*, 26(3), 479-497.

Delavande, A., Giné, X., & McKenzie, D. (2011b). Measuring subjective expectations in developing countries: A critical review and new evidence. *Journal of Development Economics*, 94(2), 151-163.

Dhami, S. (2016). *The Foundations of Behavioral Economic Analysis*. Oxford University Press, Oxford, UK.

Engelberg, J., Manski, C. F., & Williams, J. (2011). Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition. *Journal of Applied Econometrics*, 26(7), 1059-1078.

Engelmann, D., & Strobel, M. (2012). Deconstruction and reconstruction of an anomaly. *Games and Economic Behavior*, 76(2), 678-689.

Ellingsen, T., Johannesson, M., Tjøtta, S., & Torsvik, G. (2010). Testing guilt aversion. *Games and Economic Behavior*, 68(1), 95-107.

Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology*, 87(3), 327.

Epley, N., & Gilovich, T. (2016). The mechanics of motivated reasoning. *The Journal of Economic Perspectives*, 30(3), 133-140.

Eyster, E. (2002). Rationalizing the past: A taste for consistency. *Working paper*, accessed 2017/06/06, http://www.lse.ac.uk/economics/people/facultyPersonalPages/facultyFiles/ErikEyster/RationalisingThePastATasteForConsistency.pdf

Falk, A., & Zimmermann, F. (2013). A taste for consistency and survey response behavior. CESifo Economic Studies, 59(1), 181-193.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817-868.

Festinger, L. (1957). A theory of cognitive dissonance. Stanford, CA: Stanford University Press

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

Fischhoff, B. (1975). Hindsight $\neq$ foresight: the effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288-299.

Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MIT press.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. New York, Cambridge university press.

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

Guilbault, R. L., Bryant, F. B., Brockway, J. H., & Posavac, E. J. (2004). A meta-analysis of research on hindsight bias. *Basic and Applied Social Psychology*, 26(2-3), 103-117.

Guiso, L., & Parigi, G. (1999). Investment and demand uncertainty. *The Quarterly Journal of Economics*, 114(1), 185-227.

Harrison, G. W., Martínez-Correa, J., & Swarthout, J. T. (2014). Eliciting subjective probabilities with binary lotteries. *Journal of Economic Behavior & Organization*, 101, 128-140.

Harris, A. J., & Hahn, U. (2011). Unrealistic optimism about future life events: a cautionary note. *Psychological Review*, 118(1), 135.

Helweg-Larsen, M., & Shepperd, J. A. (2001). Do moderators of the optimistic bias affect personal or target risk estimates? A review of the literature. *Personality and Social Psychology Review*, 5(1), 74-95.

Hossain, T., & Okui, R. (2013). The binarized scoring rule. *The Review of Economic Studies*, 80(3), 984-1001.

Hollard, G., Massoni, S., & Vergnaud, J. C. (2016). In search of good probability assessors: an experimental comparison of elicitation rules for confidence judgments. *Theory and Decision*, 80(3), 363-387.

Holt, C. A., & Smith, A. M. (2016). Belief Elicitation with a Synchronized Lottery Choice Menu That Is Invariant to Risk Attitudes. *American Economic Journal: Microeconomics*, 8(1), 110-139

Hyndman, K. B., Terracol, A., & Vaksmann, J. (2013). Beliefs and (in) stability in normal-form games. *Working paper*, accessed 2017/06/14, http://lemma.u-paris2.fr/sites/default/files/concoursMCF/Vaksman.pdf.

Hyndman, K., Ozbay, E. Y., Schotter, A., & Ehrblatt, W. Z. (2012). Convergence: an experimental study of teaching and learning in repeated games. *Journal of the European Economic Association*, 10(3), 573-604.

Iriberri, N., & Rey-Biel, P. (2013). Elicited beliefs and social information in modified dictator games: What do dictators believe other dictators do? *Quantitative Economics*, 4(3), 515-547.

Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2), 603-606.

Khwaja, A., Sloan, F., & Salm, M. (2006). Evidence on preferences and subjective beliefs of risk takers: The case of smokers. *International Journal of Industrial Organization*, 24(4), 667-682.

Krizan, Z., & Windschitl, P. D. (2007). The influence of outcome desirability on optimism. *Psychological Bulletin*, 133(1), 95.

Krueger, J. I. (2007). From social projection to social behaviour. *European Review of Social Psychology*, 18, 1-35.

Krueger, J. I. (2013). Social projection as a source of cooperation. *Current Directions in Psychological Science*, 22(4), 289-294.

Larwood, L., & Whittaker, W. (1977). Managerial myopia: Self-serving biases in organizational planning. *Journal of Applied Psychology*, 62(2), 194

Madarász, K. (2012). Information projection: Model and applications. *The Review of Economic Studies*, 79(3), 961–985.

Manski, C. F. (2002). Identification of decision rules in experiments on simple games of proposal and response. *European Economic Review*, 46(4), 880-891.

Manski, C. F., & Neri, C. (2013) First- and second-order subjective expectations in strategic decision-making: Experimental evidence. *Games and Economic Behavior*, 81, 232-254.

Marks, G., & Miller, N. (1987). Ten years of research on the false consensus effect: An empirical and theoretical review. *Psychological Bulletin*, 102(1), 72.

McKelvey, R. D., & Page, T. (1990). Public and private information: An experimental study of information pooling. *Econometrica*, 58, 1321-1339.

Mullen, B., Atkins, J. L., Champion, D. S., Edwards, C., Hardy, D., Story, J. E., & Vanderklok, M. (1985). The false consensus effect: A meta-analysis of 115 hypothesis tests. *Journal of Experimental Social Psychology*, 21(3), 262-283.

Molnár, A., & Heintz, C. (2016). Beliefs About People's Prosociality Eliciting predictions in dictator games. *Working Paper*, accessed 2017/09/06, http://publications.ceu.edu/sites/default/files/publications/molnar-heintz-beliefs-about-prosociality.pdf

Nyarko, Y., & Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3), 971-1005.

Palfrey, T. R., & Wang, S. W. (2009). On eliciting beliefs in strategic games. *Journal of Economic Behavior & Organization*, 71(2), 98-109.

Proto, E., & Sgroi, D. (2017). Biased beliefs and imperfect information. *Journal of Economic Behavior & Organization*, 136, 186-202.

Rey-Biel, P. (2009) Equilibrium play and best response to (stated) beliefs in normal form games, *Games and Economic Behavior*, 65(2), 572-585.

Ross, L., Greene, D., & House, P. (1977). The "false consensus effect": An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13(3), 279-301.

Rubinstein, A., & Salant, Y. (2015). "Isn't everyone like me?": On the presence of self-similarity in strategic interactions. Working paper version of Rubinstein & Salant (2016), accessed 2017/09/06, https://pdfs.semanticscholar.org/34ee/9a1799fcb4c43207136437e3a1e3c3ef25a6.pdf

Rubinstein, A., & Salant, Y. (2016). "Isn't everyone like me?": On the presence of self-similarity in strategic interactions. *Judgment and Decision Making*, 11(2), 168.

Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. *Economica*, 5(17), 61-71.

Savage, L. J. (1954) The Foundations of Statistics. New York: John Wiley and Sons. (Second ed., Dover, 1972).

Selten, R., & Ockenfels, A. (1998). An experimental solidarity game. *Journal of Economic Behavior & Organization*, 34(4), 517-539.

Schlag, K. H., Tremewan, J., & Van der Weele, J. J. (2015). A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics*, 18(3), 457-490.

Schotter, A., & Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, 6(1), 103-128.

Shah, P., Harris, A. J., Bird, G., Catmur, C., & Hahn, U. (2016). A pessimistic view of optimistic belief updating. *Cognitive Psychology*, 90, 71-127.

Sutter, M., Czermak, S., & Feri, F. (2013). Strategic sophistication of individuals and teams. Experimental evidence. *European Economic Review*, 64, 395-410.

Svenson, O. (1981). Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica*, 47(2), 143-148.

Trautmann, S. T., & van de Kuilen, G. (2015). Belief elicitation: A horse race among truth serums. *The Economic Journal*, 125(589), 2116-2135.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207-232.

Tversky, A., & Kahneman, D. (1974). Heuristics and biases: Judgment under uncertainty. *Science*, 185, 1124-1130.

Van Der Heijden, E., Nelissen, J., & Potters, J. (2007). Opinions on the tax deductibility of mortgages and the consensus effect. *De Economist*, 155(2), 141-159.

Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5), 806.

Weinstein, N. D. (1989). Effects of personal experience on self-protective behavior. *Psychological Bulletin*, 105(1), 31.

Wolff, I. (2015). When best-replies are not in equilibrium: understanding cooperative behaviour. *Working paper*,accessed 2017/09/06, http://kops.uni-konstanz.de/handle/123456789/33027

Wolff, I. (2017). Lucky Numbers in Simple Games. *Mimeo*.

Yariv, L. (2005). I'll See It When I Believe It–A Simple Model of Cognitive Consistency. *Working Paper*, accessed 2017/06/06, http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.207.2893

Appendix

2.A Figures & Tables

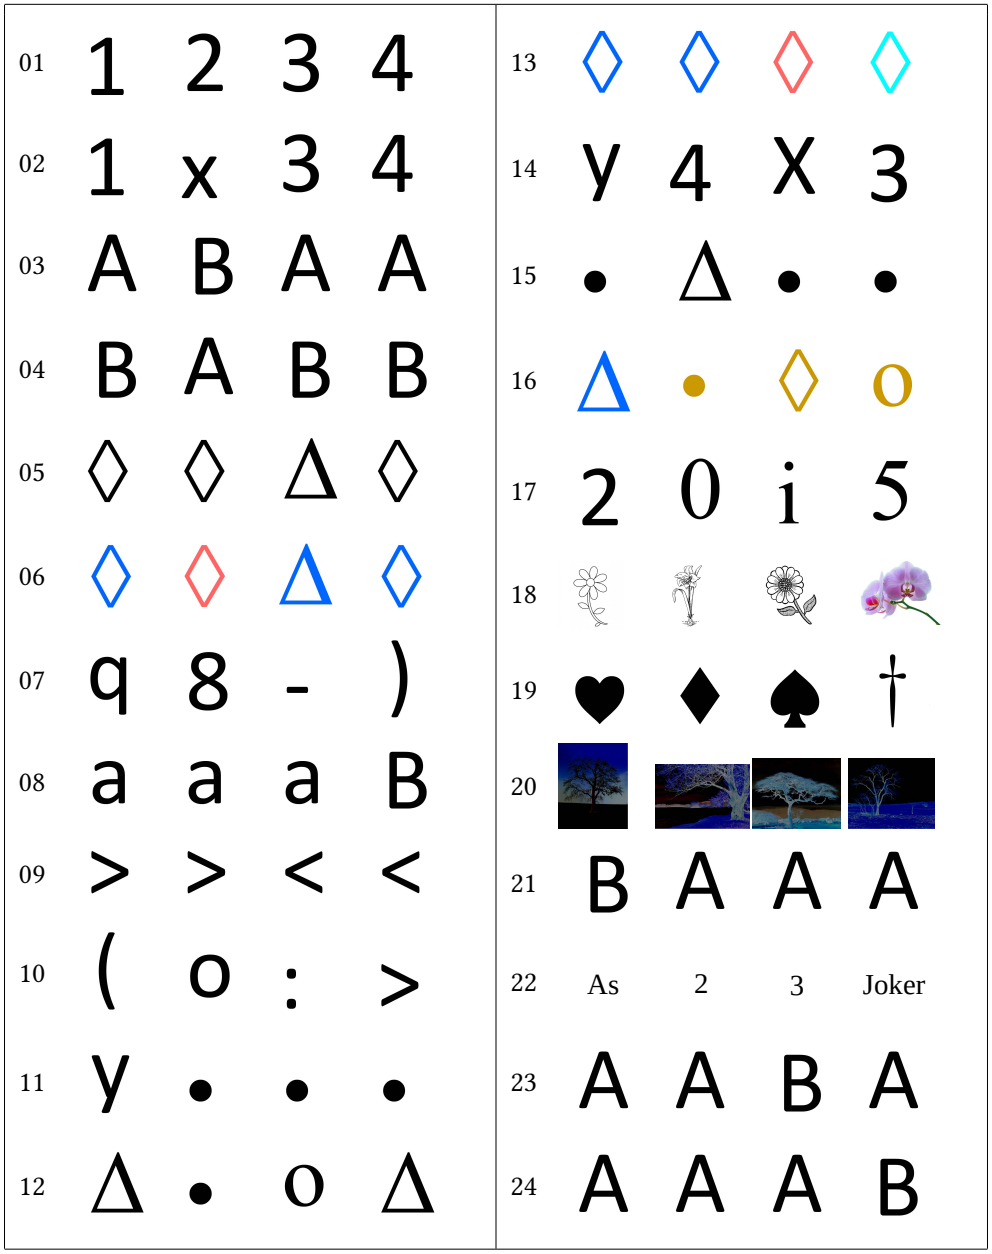

Figure A1: The 24 label sets, used to label the four options of the game. One set for each period. Sources of the pictures from sets 18 and 20 can be found in the picture credits.

Single Belief	Model $1'$
False consensus	-0.127 (2.133)
False consensus × Frame	7.677*** (2.804)
Belief to the right (GT/EPR & WT)	19.353*** (3.439)
Belief to the right (GT/EPR & WT) × Frame	-6.650* (3.926)
Hindsight Bias	-1.729 (1.820)
Hindsight Bias× Frame	1.481 (2.071)
Same Choice by the Computer	0.610 (1.121)
False consensus × Same Choice by the Computer	2.171 (2.233)
False consensus × Frame× Same Choice by the Computer	-3.127 (2.699)
Belief to the right (GT/EPR & WT) × Same Choice by the Computer	-3.787 (3.480)
Belief to the right (GT/EPR & WT) × Frame× Same Choice by the Computer	1.036 (4.077)
Hindsight Bias× Same Choice by the Computer	-0.200 (2.620)
Hindsight Bias× Frame× Same Choice by the Computer	0.983 (3.152)
Constant	20.301*** (1.032)
R^2	0.1190

Table A1: OLS dummy regressions of single belief elements with interactions for trials in which the computer (by chance) selected the same action as the participant. Standard errors in parenthesis clustered on subject level (70 clusters). Asterisks: *** $p < 0.01$, * $p < 0.1$

Single Belief	Model $1''$
False consensus	-0.251 (2.136)
False consensus × Frame	7.330*** (2.389)
Hindsight Bias	-1.810 (2.042)
Hindsight Bias × Frame	0.510 (2.017)
Belief to the right	18.448*** (2.506)
Belief to the right × Frame	-5.433* (3.104)
Constant	20.588*** (0.919)
R^2	0.1445

Table A2: OLS dummy regressions of single belief elements, used to correct beliefs. Standard errors in parenthesis clustered on subject level (70 clusters). Asterisks: *** $p < 0.01$, * $p < 0.1$

2.B Experimental Instructions

The instructions are translated from german and show the opponent frame as example. Boxes indicate consecutive screens showed to participants. The instructions of experiment 3 had the same content, but were slightly more complicated due to the belief elicitation before the action.

Today's Experiment
Today's experiment consists of 24 situations in which you will make two decisions each.

Decision 1 and Decision 2

In the first situation, you will see the instructions for bot decisions directly before the decision. In later situations, you can display the instructions again if you need to.

The payment of the experiment
In every decision you can earn points. At the end of the experiment, 2 situations are randomly drawn and payed. In one of the situations, we pay the point you earned from decision 1 and in the other situation, you earn the points from decision 2. The total amount of points you earned will be converted to EURO with the following exchange rate:

1 Point = 1 Euro

After the experiment is completed, there will be a short questionnaire. For completion of the questionnaire, you additionally receive 7 Euro. You will receive your payment at the end of the experiment in cash and privacy. No other participant will know how much money you earned.

Instructions for decision 1
In today's experiment, you will interact with other participants. **You will be randomly rematched with a new participant of today's experiment in every situation.**
Decision 1 works in the following way: You and your matching partner see th exact same screen. On the screen, you can see an arrangement of four boxes which are marked with symbols. You and the other participant choose one of the boxes, without knowing the decision of the respective other. [One of] You can earn an price of X Euro.

Experiment 1 & 3
[You only receive the X euro only if you choose **another** box than your matching partner. If both of you choose the same box, bot do not receive points in this decision]

Experiment 2
[The relative position of your chosen boxes determines who wins the price. The participant wins, whose box lies to the immediate left of the other participant's box. If one participant chooses the most left box, then the other participant wins, if he chooses the most right box. If you don't win, you receive a price of 0 euro. It is of course possible, that neither you, nor the other participant wins.]

You will only learn at the end of the experiment, which box was chosen by the other participant and which payoff you receive in a certain situation.
The arrangement of symbols on the boxes is different in every situation. Below, you can see an example of how such an arrangement could look like.

Example: The four boxes are marked from left to right by Diamond, Heart, Spade, Diamond.

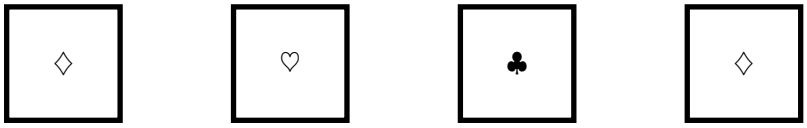

In this example, there are two boxes which are marked with the same symbol. However, the boxes on the most left and most right count as are different boxes.

Only Experiment 2

Instructions for decision 1
Although you choose a box in every situation, in some situations a box which was randomly chosen by the computer will be payoff relevant for you. This works in the following way:
After your decision, the computer draws one ball from the following urn in each situation:

If the blue ball that says "You" is drawn your own choice in decision 1 is relevant in this situation.
If the green ball that says "Computer" is drawn, the computer chooses one of the four boxes randomly (with equal probability of $\frac{1}{4}$) for you. This box will then be payoff relevant for you.
Your own decision is hence relevant with probability $\frac{1}{2}$ (=50%). The decision of the computer is relevant with probability $\frac{1}{2}$ (=50%).

The decision of your matching partner

To determine whether you won the price, we always use the original decision of your matching partner. This also holds if the computer decides for you or the other participant.
To determine whether you won the price, we hence always use the original choice of your matching partner and, depending on the drawn ball, your decision or the decision by the computer.

Text in squared brackets is frame dependent. I show the opponent frame as example.

Instructions for decision 2
In decision 2, your payoff also depends on your own decision and [on the decision of your matching partner. It will be the same matching parter, you already interacted with in decision 1.] We now explain decision 2 in detail.

Decision 2
Decision 2 refers always to a situation in which you already made decision 1. You will hence see the arrangement of boxes from the respective situation again. Again, the decision 1 [of your matching partner is relevant for you.]
Decision 2 is about your assessment, [how your matching partner decided. We are interests in your assessment of the following question:]

[*See description of frames above*]

Only Experiment 2
[Please note that decision 2 is about the **actual** (human) decision of your matching partner and **not** about a possible computer decision.]

For every box, you can report your assessment [with what probability your matching partner chose the respective box]. You can enter the percentage numbers in a bar diagram. By clicking into the diagram, you can adjust the height of the bars. You can adjust as many times as you like, until you confirm.
Since your assessments are percentage numbers, the bars have to add up to 100%. The sum of your assessment is displayed on the right. You can adjust this value to 100% by clicking. Or you enter the relative sizes of your assessments only roughly and then press the "scale" button. Please note, that because of rounding, the displayed sum ma deviate from 100% in some cases.
On the next page, we explain the payoff of decision 2.

Text in squared brackets is frame dependent. I show the opponent frame as example.

The payoff in decision 2
In this decision, you can either earn 0 or 7 points. Your chance of earning 7 points increases with the precision of your assessment. Your assessment is more precise, the more it is in line with [the decision behavior of your matching partner. For example, if you reported a high assessment on the actually selected box, your chance increases. If your assessment on the selected box was low, your chance decreases.]
You may now look at a detailed explanation of the computation of your payment, which rewards the precision of your assessment.

It is important for you to know, that the chance of receiving a high payoff is maximal in expectation, if you assess the behavior of your matching partner correctly. It is our intention, that you have an incentive to think carefully about the behavior of your matching partner. We want, that you are rewarded if you have assessed the behavior well and made a respective report.

Your chance will be computed by the computer-program and displayed to you later. At the end of the experiment, one participant of today's experiment will roll a number between 1 and 100 with dies. If the rolled number is smaller or equal to your chance, you receive 7 points. If the number is larger than your chance, you receive 0 points.

Text in squared brackets is frame dependent. I show the opponent frame as example.

Payment of the assessments
At the end of your assessment, you will receive the 7 points with a certain chance (p) and with $(1-p)$, you receive 3 points. You can influence your chance p with your assessment in the following way:

As described above, you will report an assessment for each box, on how likely [your matching partner is to select that box. One of boxes is the actually selected. At the end, your assessments are compared to the actual decision of your matching partner.] Your deviation is computed in percent.

Your chance p is initially set to 1 (hence 100%). However, there will be deductions, if your assessments are wrong. The deductions in percent are first squared and then divided by two.

For example, if you place 50% on a specific box, but [your matching partner selects another box,] your deviation is equal to 50%. Hence, we deduct $0.50{*}0.50{*}\frac{1}{2} = 0.125$ (12.5%) from p.

[For the box, which is actually selected by your matching partner, it is bad if your assessment is far away from 100%. Again, your deviation from that is squared, halved and deducted. For example if you only place 60% probability on the actually selected box, we will deduct $0.40 * 0.40 * \frac{1}{2} = 0.08$ (8%) from p.]

With this procedure, we compute your deviations and deductions for all boxes.
At the end, all deductions are summed up and the smaller the sum of squared deviations is, the better was your assessment. For those who are interested, we show the mathematical formula according to which we compute the quality of your assessment and hence your chance p of receiving 7 points.

$$p = 1 - \frac{1}{2}\left[\sum_i (q_{box_i,estimate} - q_{box_i,true})^2\right]$$

The value of p of your assessment will be computed and displayed to you at the end of the experiment. The higher p is, the better your assessment was and the higher your chance to receive 7 points (instead of 0) in this part. At the end of the experiment, the computer will roll a random number between 0 and 100 with dies. If this number is smaller or equal to p, you receive 7 points. If the number is larger than p you receive 0 points.

Summary
In order to have a high chance to receive the large payment, it is your aim to achieve as few deductions from p as possible. This works best, if you have an good assessment of the behavior of participant B and report that assessment truthfully.

Chapter 3

Belief Uncertainty and Stochastic Choice[§]

Dominik Bauer Irenaeus Wolff

Graduate School of Decision Sciences & Thurgau Institute of Economics

University of Konstanz, Germany

Abstract: People often cannot assign a clear probability to an event but face uncertainty about their subjective probabilities. We model belief uncertainty by assuming that agents' beliefs are characterized by a distribution over subjective-probability distributions that agents cannot access directly. Our model produces stochastic choice because each decision-relevant belief is but one realization out of the distribution over all possible beliefs. Our model predicts that when comparing unknown situations to routine choices, people will make more *ex-ante* suboptimal choices in unknown situations. The model also offers an explanation for experiment participants not playing a best-response to their stated beliefs: participants are uncertain which belief to report or base their decision on, and hence, act on momentaneous 'belief realizations'. In an experiment, we manipulate participants' belief uncertainty. We do so by exogenously manipulating their strategic uncertainty, providing varying levels of information about historical choice data. In situations in which belief uncertainty is low, observed best response rates are high and increasing in the amount of information we provide. Conversely, high belief uncertainty leads to lower consistency. On top, low belief uncertainty makes participants earn more, even when controlling for the accuracy of their beliefs.

JEL classification: C91, D81, D83

Keywords: Belief Elicitation, Belief-Action Consistency, Stochastic choice, Discoordination Game

[§]We thank Fabian Dvořák, Thomas Hattenbach, Georg Weizsäcker, Ian Krajbich, Wieland Müller, Tomasz Strzalecki, Wolfgang Luhan, Alexander K. Wagner, the research group at the Thurgau Institute of Economics, the members of the Graduate School of Decision Sciences of the University of Konstanz, as well as participants of the Thurgau Experimental Economics Meeting (theem) 2017 and the ESA European Meeting 2017 for helpful comments.

Contact: Chair of Applied Research in Economics, University of Konstanz, Universitätsstraße 10, D-78464 Konstanz, Germany.

3.1 Introduction

This paper is about the stochasticity of choices in decisions under uncertainty. In particular, the paper is about the likelihood of *ex-ante* suboptimal choices, and about the degree of consistency we should expect when a person has to make multiple similar or even identical decisions. We claim that the person's behavior will be more inconsistent and error-prone the higher the degree of uncertainty in her belief.

To motivate what we mean by uncertain beliefs, consider the following example: We know for sure what the odds of a fair coin flip are. When offered a bet on this coin flip, it is easy to see whether it is worth accepting the bet or whether the odds-maker tries to trick us. Now imagine a colleague offers you a bet over a bottle of wine on your favorite football team winning the next match. It is the final match for the championship, your team is the home team, and your team performed better overall during the season. But then, one of the top scorers of your team is injured. So, you start thinking about your belief on how likely your team is to win the match. You would probably say: err, let me see, probably chances are 50% that they win. But just how certain would you be that it's not 60%, or 40%, for that matter?

Both for the coin flip and the football match it would be sensible to report a fifty-fifty belief when asked for it. However, there is some difference in confidence about that statement. For the coin flip, there is no sensible answer other than fifty-fifty. For the football-match, it could be also 60% or 40%, while fifty-fifty seems like a reasonable *average* answer. Arguably, people face this kind of uncertainty about the probability of events very often. Most of the time, they would not even be able to put real numbers on the probabilities. For our purposes, however, it suffices to model the general situation in terms of compound risk, rather than in terms of uncertainty in its narrower sense.

We propose a model of stochastic choice and belief reports, where the stochasticity is due to belief uncertainty. In our model, we follow the standard Bayesian approach of assuming that agents have subjective probabilistic beliefs over the sources of uncertainty they face. In our example of which football team is going to win the match, the agent's assessment therefore is a probability distribution over probabilities. Generally, such a probability-distribution over probabilities simply defines a compound lottery. Standard models assume that agents can reduce this compound lottery to a simple lottery, and the resulting probability assessment is then interpreted as 'the belief'.

However, we assume that agents cannot directly access the probability distribution over probabilities and hence, they cannot reduce it. Rather, when acting *and* when reporting a belief,

agents will sample a probability from the probability distribution and react to the randomly drawn probabilities as if they were the true probability.[1] We call the distribution over probabilities a *belief distribution*, because it is a distribution over different possible beliefs. If the underlying belief distribution is spread out and many different beliefs are likely to be drawn, the agent is *uncertain* about what the 'true', that is, the reduced probability is. We call the variance of the belief distribution its *belief uncertainty*.[2]

Our model has important implications. It will produce stochastic choice, and agents will make errors, where we define an error as a choice inconsistent with the reduced probability under the agent's belief distribution. Moreover, the variance of an agent's expected choice in a given situation—and hence, the probability of the agent making an error—will depend on the variance of the belief distribution. This will be the case even when the reduced probability is held constant. Thus, our model relates the uncertainty of a situation to the likelihood of *ex-ante* suboptimal choices. To the best of our knowledge, this feature is unique to our model.

Relating uncertainty to suboptimal behavior is important because many of the important decisions we make in life are decisions that are rarely repeated—and therefore, characterized by a high degree of uncertainty. Uncertainty also plays a critical role, for example, in investment decisions. There is a whole literature on whether investment decisions will be affected by uncertainty, and in what way (*e.g.*, Guiso & Parigi, 1999). While the empirical literature often relates to uncertainty in its narrower sense (*e.g.*, Baker, Bloom & Davis, 2016), the theoretic arguments mostly focus on compound risk. Our study can inform this literature in that uncertainty will not only affect rational decision-making in predictable ways, but that it will also increase the prevalence of *ex-ante* erroneous investment decisions. A prediction of our model therefore would be that following a time of high uncertainty we should see a higher number of bankruptcies particularly of small firms compared to after a time of little uncertainty.

By introducing a new form of stochasticity, our model also provides a new perspective on learning in unknown situations. When facing a decision for the first time, belief uncertainty is likely to be high. In our model, this leads to high error rates. Hence, there is scope—and need—for learning. As the situation is repeated with feedback, the agent gathers more and more observations. In most situations, gathering more information will decrease the variance of the belief distribution, leading to less errors. Hence, the agent learns how to behave in the situation by identifying the situation better and better, even when there is no change in the reduced belief. This sets us apart, for example, from models of belief-based learning like fictitious play or

[1] A similar idea is the idea of discovered preferences, see Plott (1996) or Cubitt, Starmer & Sugden (2001).

[2] See Pouget, Drugowitsch, & Kepecs (2016) for a neuroscience perspective on uncertainty. Just like we do, the authors define uncertainty about some proposition as the variance of a posterior distribution (p. 369).

Cournot learning. Finally, the stochasticity due to belief uncertainty also concerns ourselves as researchers interpreting results from economic experiments. It provides a possible explanation for the widely varying belief-action-consistency rates reported in the literature.[3]

We test the implications of our model on belief-action consistency in a laboratory experiment. Our participants play a series of discoordination-games. Instead of matching them within-session, they play against one choice out of a distribution of choices from the same game, but from earlier sessions. Before they play the game and report a belief, participants receive information on the relevant distribution of choices from the earlier sessions. To manipulate belief-uncertainty, we give participants samples of varying size from the distribution they play against. Participants have to combine this information with their prior beliefs to arrive at a posterior belief distribution. The additional information can have two effects: if the information supports the participant's prior judgement of what the distribution of choices will be, more information will increase the participant's faith in her belief. If, on the other hand, the information goes against the participant's prior judgement, more information may actually disconcert the participant more. This is indeed what we see: if the provided information is congruent with the prior belief—and hence belief uncertainty decreases—the observed best-response rate is higher on average and increases in the number of observations we provide. When the information is not in line with prior beliefs, uncertainty increases and belief-action consistency decreases.
Our results are robust to controlling for the costs of making an error which determine the probability of decision errors in some classical stochastic-choice models. The results are in line with our model predictions: making participants more certain about the relevant underlying process (the object of their belief) leads to less stochasticity of actions and belief reports and, hence, to higher consistency. Moreover, participants on average earn higher payoffs when uncertainty is low, even when we control for the empirical accuracy of their beliefs. This paper thus showcases the importance of an additional—but so far, neglected—source of stochastic choice and its consequences for observed behavior in experiments.

Related Literature

Many of the ideas behind our motivation for this study can be found already in the literature on choice under uncertainty.[4] The main question in this literature is how people will make their

[3] *E.g.*, Costa-Gomez & Weizsäcker (2008) find best-response rates as low as 51% in some of their 3x3 games, while Manski & Neri (2013) report a best-response rate of 89% in a 2-action Hide-and-Seek game.

[4] Even our introductory example in Section 3.1 is similar to the examples given in this literature, *cf., e.g.*, Gilboa, Postlewaite, and Schmeidler (2008).

decisions when they face uncertainty and there is no clear way of assigning probabilities to the possible states of the world. This literature departs from Savage's (1954) idea that when agents face ambiguity, they will simply form subjective beliefs and act on those subjective beliefs as if the beliefs were proper probabilities. There is a whole array of how the corresponding non-Bayesian subjective probabilities are modelled, and how they are used by the agents.[5] The approach that is probably closest to ours is the multiple-priors approach axiomatized in Gilboa & Schmeidler (1989). In multiple-priors models, agents generally choose among the alternatives using a maximin-utility criterion across all probabilities they consider possible.

The literature on choice under uncertainty nicely explains ambiguity aversion as exemplified by the Ellsberg paradox. It also explains, for example, the fact that most people will neither buy a stock nor sell it short for a whole price range, rather than being indifferent between either of these options and not doing anything only at one specific single price.[6] So, generally speaking, this literature focuses on explaining choices. Our aim complements this literature, as we focus on explaining the variance within people's choices, and on the likelihood of observing inconsistent choices and errors. For this purpose, it is sufficient to slightly adapt the Bayesian model. This does not mean that we are convinced that people in reality can always come up with a probability distribution nor handle compound lotteries or update such a distribution in a Bayesian way. We use our modified Bayesian model merely as a tractable as-if description.

There is also a huge and important literature on stochastic choice. This literature started from the attempt to explain effects like the Allais paradoxes. In principle, our model is mute on these effects because in any of the Allais paradoxes, the probabilities are given, and therefore, obvious to the decision-maker. However, we think that our model still provides a possible intuition for these cases, namely, if we consider some of the options to be too complex for the agent to assess them directly (*e.g.*, because the agent is not good at handling probabilities). In this case, the agent may have to form a subjective belief about how good each option is. This subjective belief may then be more uncertain, the more complex the option is. This can in principle give rise to the commonly observed effects in Allais-type tasks analogously to the way Fechner-type errors do. We discuss the specific modelling differences to the common stochastic-choice models in Section 3.2.1.

Last but not least, our paper is related to the vast literature on learning. However, the only type of learning that has been documented in the literature and that could interfere with our

[5] For a recent review, *cf.* Etner, Jeleva, and Tallon (2012), who also discuss important economic applications of the models.

[6] The first solution to the buying-a-stock problem is due to Dow & Werlang (1992), using Schmeidler's (1989) Choquet expected utility.

conclusions is feedback-less learning (Weber, 2003). According to this idea, experiment participants learn how to play a game even without feedback. We therefore should see increasing best-response rates over time. We address this potential confound by (individually) randomizing the order in which participants receive the different sample sizes. On top, we control for the period in our analysis, and hence, implicitly also for any form of feedback-less learning on how to play a best-response.[7]

3.2 A model of belief uncertainty and stochastic choice

In this section, we present the simple example of a two-player two-action discoordination game to make our point. Of course, our model applies also to more general settings. First, we present our model of belief uncertainty and contrast it with other stochastic choice models from the literature. Then, we relate it to observed best-response rates and present consequences of information updating for error rates at the end.

Our model is a model of individual choice. The model is not a game-theoretic model even though in our main example, the object of agents' beliefs is the behavior of the other player. While it would be conceivable in principle to extend the model to an equilibrium model akin to a quantal-response equilibrium, this is not the focus of our study. For our main example, it is even essential that agents do not hold equilibrium beliefs. Non-equilibrium beliefs are essential because with equilibrium beliefs there cannot be any errors in a pure discoordination game.[8] It also would be possible to incorporate our model into a cognitive-hierarchy model, but our point here merely is to highlight the importance of uncertainty in people's beliefs that goes beyond the effect of the cost of making an error. In this sense, we abstract from the question of where people's initial beliefs and the initial belief uncertainty come from.

[7] There are (at least) three additional broad categories of learning models. Directional learning (Selten and Stöcker, 1986) does not apply because it is tailored to situations were the decision variable is on at least an ordinal scale. Experience-weighted attraction learning (Camerer & Ho, 1999) and belief-based learning (such as fictitious play, Brown, 1951, or Cournot play) can be interpreted in a way that makes them applicable to our setting. In that case, the two models would make the prediction that participants could be learning something from the information we provide. However, the predictions in this case are hardly distinguishable from the predictions of (noisy) standard theory. Under belief-based learning, participants could be trying to learn the mixed strategy of their opponent from the information, assuming a homogeneous population. In that case, they simply should be best-responding to the information we provide. Even if agents were to learn more broadly how others behave 'in this type of situation', they still should be playing a best-response to their beliefs. Under experience-weighted attraction learning, participants could update their initial choice propensities using the information we provide, again assuming a homogeneous population. The resulting behavior should be very similar to best-responses with Fechner-type errors.

[8] The assumption of non-equilibrium beliefs seems warranted given the experimental evidence, *e.g.*, for four-action discoordination games. In the data from Bauer & Wolff (2017) we use also for our present experiment, choice distributions are significantly different from uniformity at a 5%-level in 15 out of 24 settings (χ^2-test). 15 out of 24 settings are clearly more than the expected 1.2 settings under equilibrium behavior. See Table C1 for the data.

3.2.1 Model

An agent plays a simultaneous two-player discoordination game with two options, X and Y. She is randomly matched with another player out of a population of N players. If she chooses a different action than the other player, both receive a payoff of $\pi = 1$, and nothing, otherwise. Assume players have any commonly used utility function, potentially displaying non-neutral risk- and loss-attitudes or being driven by social preferences. In all of these cases, the best-response in the discoordination game depends only on the probability $\hat{\phi}$ the player assigns to the other player choosing X:

$$BR(\hat{\phi}) = \begin{cases} X, & \text{if } \hat{\phi} < 0.5 \\ (X, p;\ Y, 1-p) \ \mid \ p \in [0,1] & \text{if } \hat{\phi} = 0.5 \\ Y, & \text{otherwise.} \end{cases} \tag{3.1}$$

The true probability ϕ^* of X-choices in the population is an unknown realization of the random variable $\Phi \in [0,1]$. Mean and variance of Φ are unknown as well. Hence, the player has to rely on a belief about ϕ. In this paper, we assume that the belief is a non-degenerate probability distribution over all possible values of ϕ. For example, the player might assign a probability of 40% to $\phi = 0.7$ and distribute the remaining 60% of the probability mass over all other possible values of ϕ. Hence, the belief is a probability distribution $\phi \sim (\mu_q, \sigma_q)$ with continuous density function $q(\phi)$ where $\int_0^1 q(\phi)d\phi = 1$ and $q(\phi) > 0, \forall \phi$. Considering this belief distribution, the player faces a compound lottery: with density $q(\phi')$ the other player chooses X with probability ϕ'. However, in standard theory this subtlety does not play a role, as the best-response depends only on the *expected* probability the agent assigns to the other player choosing X, denoted by:

$$E_q[\phi] = \int_0^1 \phi \cdot q(\phi)d\phi = \mu_q \tag{3.2}$$

In standard theory, the player will then choose $BR(\mu_q)$. In the two-option case we outline here, the critical belief in (3.1) happens to be $\phi^{crit} = 0.5$. This need not be the case in games with different payoffs or more than two options. Also note that, as outlined above, we ignore the (trivial) case of equilibrium beliefs because in equilibrium, any action always is a best-response. Throughout the whole paper, we will hence assume $\mu_q \neq 0.5$.

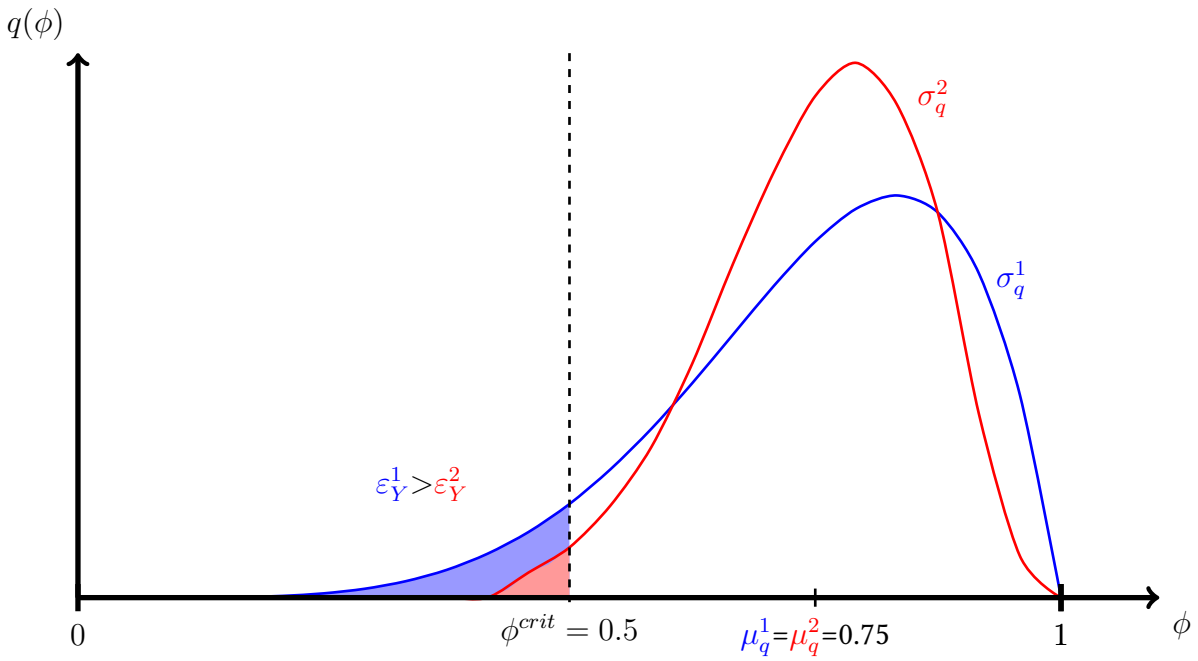

Figure 1: Two belief distributions with identical means but differing variances. The shaded areas indicate the respective error rate ε_Y

Stochastic choice

We propose that players do not have direct access to $q(\phi)$ when playing a game or reporting beliefs. Therefore, they cannot compute μ_q. Instead we assume that whenever the belief is consulted, the player draws one value ϕ^r from $q(\phi)$. This draw ϕ^r is then used to determine the optimal action $BR(\phi^r)$ instead of $BR(\mu_q)$. Hence, not only the mean but also the whole distribution $q(\phi)$ matter for players' predicted choices.

In contrast to standard theory, players will make errors in our model. We define the error rate as the probability that the player draws a ϕ^r that does not indicate the same best response as μ_q, that is, the probability that $BR(\mu_q) \neq BR(\phi^r)$. Consider the example distributions in Figure 1 with $\mu_q > 0.5$ so that $BR(\mu_q) = Y$. Then the error rate is characterized by the probability mass of $q(\phi)$ on all $\phi < 0.5$ and indicated by the shaded areas. Denote by ε_k with $k \in \{X, Y\}$ the error rate conditional on $BR(\mu_q)$. Define $\varepsilon_Y = \int_0^{0.5} q(\phi)d\phi = Q(0.5)$ as the error rate in case $BR(\mu_q) = Y$. In case $BR(\mu_q) = X$, the error rate is $\varepsilon_X = \int_{0.5}^1 q(\phi)d\phi = 1 - Q(0.5)$.

Errors in our and other stochastic choice models

The widely used models of stochastic choice do not use belief uncertainty as a direct source of stochasticity. Tremble-error models (Harless & Camerer, 1994) assume a constant error when executing a decision. Our idea of stochastic choice is in a way more related to random-

preference models (Becker, DeGroot & Marschak, 1963; Loomes & Sugden, 1995). However, instead of a probability distribution over parameters of the utility function, we assume a distribution over beliefs. Also, in our specific example of the discoordination game, random-preference-models do not predict errors, because optimal behavior is invariant to changes of the utility-function in this game. In our model, two characteristics of the belief distribution determine the error rate.

First, as the mean μ_q approaches 0.5, the error rate ε_k increases *ceteris paribus*. The closer the belief is to indifference, the more errors are made. This is also predicted by models in which errors depend on expected-utility differences (ΔEU), like Fechner-error (Becker, DeGroot & Marschak, 1963; Fechner 1860/1966; Hey & Orme, 1994), Quantal-Response-Equilibrium (McKelvey & Palfrey, 1995) or Drift-Diffusion models (Ratcliff, 1978). However, in these models, errors happen for different reasons. Applying such a model to our setting, the probability of committing an error depends directly on the distance of the belief mean to indifference: $\Delta EU = |\mu_q - 0.5|$ because of external random shocks which disturb the original utility. In our model, the error rate is *endogenously* determined by the belief distribution and depends on both ΔEU *and* the probability dispersion around the mean. As the other models, it predicts increasing error rates as ΔEU approaches zero, but due to the shift of probability mass across the critical threshold and hence depending on the variance of $q(\phi)$. In our model, it is for example possible that a belief with ΔEU close to zero produces little or no errors if the variance of $q(\phi)$ approaches zero. Our model therefore provides an intuition for when people will violate the monotonicity principle and choose stochastically dominated options. Violations of monotonicity are one of the greatest challenges for stochastic-choice models: people often violate monotonicity when dominance is not obvious (because their belief over which option is better is uncertain). On the other hand, people respect monotonicity when dominance is obvious (and therefore, they know the best option exactly).

Second, for the error rate ε_k to increase, in our model it is sufficient that *ceteris paribus* the variance of $q(\phi)$ increases. Consider again Figure 1. The shaded areas are the values of ε_k for two belief distributions with the same mean ($\mu_q^1 = \mu_q^2$) and hence $\Delta EU^1 = \Delta EU^2$, but different variances ($\sigma_q^1 \neq \sigma_q^2$). The more variance $q(\phi)$ has around its mean, the more likely the agent commits an error. When drawing from the high-variance belief, it is more likely that $BR(\phi^r) \neq BR(\mu_q)$ compared to a draw from the low-variance belief. Our model endogenizes the error rate and predicts that higher belief uncertainty leads to more errors. This separates our model from other stochastic choice models and to the best of our knowledge, there is no other model that relates additional characteristics of $q(\phi)$—like the variance—to the probability

of an error. Also, we do not rule out other sources of error: after treating the draw of ϕ^r as the "true" belief, any of the other models of error may apply. Put differently, our model can be applied *on top* of the other models.

Stochastic beliefs

The notion of stochastic choice has consequences also for belief reports. In the usual experiment, choosing an action and reporting a belief are two separate decisions with different incentives. The reported beliefs are usually assumed to approximate μ_q and used to explain behavior. They are interpreted as the true cause of an action. We relax this interpretation by assuming that not only the actions but also the belief reports are stochastic. Instead of calculating and reporting μ_q as a belief, the player also reports one draw ϕ^r as a belief. We assume that players use two different and independent draws from $q(\phi)$ for the two tasks.[9] Denote by ϕ^r_A the draw used for the action and by ϕ^r_B the draw for the belief report. Below, we will discuss the consequences of the combination of stochastic choice and stochastic belief reports for consistency. Note, however, that for our general predictions it would be sufficient to assume that either the action-relevant belief *or* the belief for the report are drawn randomly (while maintaining the standard assumptions of a best-response to μ_q, or a truthful report of μ_q, respectively). We nonetheless assume both belief draws to be stochastic. On the one hand, a stochastic ϕ^r_A makes our theory applicable also to individual-choice settings and makes it easily comparable to existing models of stochastic choice. On the other hand, non-stochastic belief reports seem implausible once we assume stochastic choices due to stochasticity in beliefs.

So far, we have introduced the key idea that when making decisions and when reporting beliefs, agents have to draw realizations from their inner belief distribution. We have characterized the error rate and we have contrasted this rate to what would be predicted by other existing models of stochastic choice. We now turn to the implications of our model for observed behavior in experiments.

3.2.2 Belief-action consistency

We assume both choices and belief reports to be stochastic. Hence, the true belief distribution $q(\phi)$ and therefore also the true best-response rate and the true error rate are usually unobservable in experiments.

For the experimenter to *observe* consistent behavior, that is, an action that is a best-response

[9] If a single draw were to determine both action and belief, we would predict a 100% best-response rate which definitely is not what we observe.

to the reported belief, the two draws from the belief distribution have to 'fit together'. A best-response is observed only if $BR(\phi_A^r) = BR(\phi_B^r)$. In our example above, this is the case whenever both $\phi_A^r, \phi_B^r > 0.5$ or both $\phi_A^r, \phi_B^r < 0.5$. The expected observed best-response rate $\widehat{BR}$ is directly connected to the error rate ε_k defined earlier and can be characterized by:

$$\widehat{BR} = Prob\Big[BR(\phi_A^r) = BR(\phi_B^r)\Big] = \varepsilon_k^2 + (1-\varepsilon_k)^2 \tag{3.3}$$

A best response is observed if an error occurs in either none or both of the draws ϕ_A^r, ϕ_B^r. To obtain further results, we need to put some structure on the belief distribution $q(\phi)$. We assume $q(\phi)$ to be a Beta-distribution which is a very flexible distribution that is able to approximate many different distributions over beliefs in our setting.[10]

PROPOSITION 1: If $q(\phi; \alpha, \beta)$ with $\mu_q \neq 0.5$ is the Beta-distribution with hyperparameters $\alpha, \beta > 1$, the expected observed best-response rate $\widehat{BR}$ decreases in the error rate ε_k in a symmetric game.

PROOF: $\frac{\partial \widehat{BR}}{\partial \varepsilon_k} = 4\varepsilon_k - 2$. Hence, $\widehat{BR}$ decreases in ε_k if $\varepsilon_k < 0.5$. The error rate ε_k is always smaller than 0.5, if the median m_q of the belief distribution $q(\phi; \alpha, \beta)$ is on the same side of the critical value as the mean μ_q (that is, if the median favors the same best response as the mean $BR[m_q] = BR[\mu_q]$) because then, more than 50% of the probability mass are contained in $(1-\varepsilon_k)$. For the symmetric games we consider here, it is hence sufficient to show that either **both** or **neither** the mean and median of $q(\phi)$ are larger than the critical value of $\phi^{crit} = 0.5$.
By the mode-median-mean inequality (Groeneveld & Meeden, 1977), $\mu_q \leq m_q$ if $1 < \beta < \alpha$. However, if $\beta < \alpha$, also $\mu_q = \frac{\alpha}{\alpha+\beta} > 0.5$. Hence, if $1 < \beta < \alpha$, then $0.5 < \mu_q \leq m_q$ (and analogously, $m_q \leq \mu_q < 0.5$ if $1 < \alpha < \beta$). ■

Note that PROPOSITION 1 also holds if either the action *or* the belief are assumed to be non-stochastic. In these cases, the expected observed best response rate is simply $\widehat{BR}' = (1-\varepsilon_k)$ and obviously $\frac{\partial \widehat{BR}'}{\partial \varepsilon_k} < 0$.
Having specified how the observed belief-action consistency in experiments will depend on belief uncertainty, we next look at a possible determinant of belief uncertainty. A natural source

[10] The Beta-distribution is a prominent example of a probability density function with support (0,1) and hence suitable to model a distribution over probabilities. With this distributional assumption, it will be convenient to apply Bayesian-updating, as the Beta-distribution is a conjugate prior for the Bernoulli and Binomial distributions. Hence, updating a prior belief (Beta-distributed) by a number of X-choices in a sample (n i.i.d. Bernoulli variables) will again yield a Beta-distributed posterior. See section 3.2.3.

of variation in the belief distribution—and hence also in belief uncertainty—is the integration of new information into the belief. To pave the ground for the hypotheses for our experiment, we will explore the influence of information integration on the error rate in the following section.

3.2.3 Bayesian updating

From now on, let $q(\phi; \alpha, \beta)$ denote the participant's prior belief distribution. The mean of the Beta-distribution and hence the prior mean is $\mu_q = \frac{\alpha}{\alpha+\beta}$. The hyperparameter $\alpha = n_X^{Prior} + 1$ can be interpreted as the number of prior observations of X-choices in a sample of $n^{Prior} = n_X^{Prior} + n_Y^{Prior}$ choices and $\beta = n_Y^{Prior} + 1$ as the number of prior observations of Y-choices. Suppose the player observes a new sample of $n = n_X + n_Y$ decisions from the population of N other players, where n_X denotes the number of X-choices in the sample. The likelihood function of ϕ, given the observed sample is $s(\phi|n) = \phi^{n_X} \cdot (1-\phi)^{n_Y}$. The sample mean is defined as the share of X-choices in the sample $\mu_s = \frac{n_X}{n_X+n_Y}$. The player updates her prior belief about ϕ according to Bayes' rule to obtain the posterior $p(\cdot)$ with mean μ_p:

$$p(\phi|n, \alpha, \beta) = \frac{s(\phi|n) \cdot q(\phi; \alpha, \beta)}{t(n, \alpha, \beta)} \tag{3.4}$$

where $t(n, \alpha, \beta) = \int_0^1 s(\phi|n) q(\phi; \alpha, \beta) d\phi$. Because of conjugacy, the posterior is Beta-distributed as well. Hence now $\phi \sim Beta(\alpha + n_X, \beta + n_Y)$.

Posterior mean and variance

The posterior mean can be written as:

$$\mu_p = \frac{\alpha + n_X}{(\alpha + n_X) + (\beta + n_Y)} = \underbrace{\frac{\alpha+\beta}{\alpha+\beta+n}}_{1-w} \cdot \underbrace{\frac{\alpha}{\alpha+\beta}}_{\mu_q} + \underbrace{\frac{n}{\alpha+\beta+n}}_{w} \cdot \underbrace{\frac{n_X}{n}}_{\mu_s} \tag{3.5}$$

The posterior mean is hence a weighted combination of the sample- and the prior-mean. The weights are determined by the relative number of observations in the respective distribution where w denotes the relative weight of the sample. Further note that $\lim_{n\to\infty} \mu_p = \mu_s$.
The posterior's variance can be expressed as $\sigma_p = \frac{\mu_p(1-\mu_p)}{\alpha+\beta+n+1}$. It has two important properties. First, as $\frac{\partial \sigma_p}{\partial n} < 0$ the variance decreases *ceteris paribus* in n, the number of observations in the sample. Second, the variance is inverse U-shaped with a maximum at indifference $\mu_p = 0.5$. Hence, the variance decreases *ceteris paribus* in the distance of the belief mean to indifference $|\mu_p - 0.5|$.

The error rate of the posterior

As described above, the sample- and prior means as well as their relative weight determine the location and shape of the posterior belief distribution. In this section we derive predictions for the posterior's error rate ε_k based on characteristics of the prior and the observed sample. In the following, we continue to assume $BR(\mu_q) = Y$ for simplicity, but all predictions hold symmetrically for priors with $BR(\mu_q) = X$. The most important characteristic is the location of μ_s relative to μ_q and to the critical threshold from equation (3.1), in our case, to 0.5. There are three cases:

I) Congruent sample: The sample mean is the same or greater than the prior mean: $0.5 < \mu_q \leq \mu_s$.

i) If $0.5 < \mu_q < \mu_s$ then ε_k decreases as the posterior mean is shifted to the right and hence, probability mass is shifted away from 0.5.

ii) If $0.5 < \mu_q = \mu_s$ then ε_k decreases as the posterior variance decreases.

In both of these subcases, an increase of the relative weight of the sample w leads to an additional decrease of posterior variance and, hence, a larger decrease of ε_k.

II) Sample in between: The sample mean is less extreme than the prior but favors the same action: $0.5 < \mu_s < \mu_q$. In this case, the prediction depends on the relative weight.

i) For a sufficiently small relative weight, ε_k will increase due to the shift of the mean towards 0.5 which is stronger than the minor decrease of variance.

ii) For a sufficiently large relative weight of the sample, ε_k will decrease as the decrease of variance of the posterior will outweigh the effect of the shift towards 0.5.

III) Incongruent sample: If $\mu_s < 0.5 < \mu_q$, that means, if the sample mean is completely different from the prior mean and the two suggest different best-responses, it is *a priori* unclear which action the posterior will favor. The prediction depends again on the relative weight:

i) For a sufficiently small relative weight of the sample, ε_k increases as long as the posterior mean μ_p is such that $\mu_s < 0.5 \leq \mu_p < \mu_q$. This means, the posterior mean approaches 0.5 from the right and probability mass is shifted to the left.

ii) If the relative weight is large enough, the prior is 'overturned' by the information. Then, μ_s outweighs μ_q and $\mu_p < 0.5$. From then on, ε_k decreases in relative weight (from $\varepsilon_k^{max} = 0.5$ at $\mu_p = \lim_{\epsilon \to 0} 0.5 - \epsilon$).

Note that in Cases I and II, the posterior will always favor the same action as the sample because both $\mu_q, \mu_s > 0.5$. This also holds for the overturned beliefs in case III ii). However, if the belief is not overturned, the posterior will always favor a different action than the sample in case III i).

In this section, we have shown how the integration of new information affects belief uncertainty and the error rate. In Section 3.4, we will use the outlined cases to formulate specific hypotheses for our experiment, which we describe next.

3.3 Experiment

We test the predictions of our model on the relationship of belief uncertainty and belief-action consistency in an experiment. We manipulate belief-uncertainty exogenously by giving varying amounts of information about the decisions of the relevant target population of other players.

Experimental tasks

The experiment uses a two player, four-option, one-shot discoordination-game. Participants play a series of 24 games without any feedback in between and are randomly rematched before every game. The four options of each game are labeled boxes. If a participant chooses another box than her current matching partner, both receive 7€ and nothing otherwise. The labels of the four options vary in every game and we use a large variety of letters, numbers or symbols as labels. For example, we start with labels "1,2,3,4" in game 1 and "1,x,3,4" in game 2. Hence, only the non-strategic features of the game vary across periods. The complete list of all labels is depicted in Figure C1 in the appendix. The order of the games is the same for all participants. Along with every choice in the game, we elicit probabilistic beliefs after the action for every period. Participants have to report a set of four probabilites, one for each box. We incentivise the belief reports via a Binarized-Scoring Rule (Hossain & Okui 2013) where subjects could earn another 7€. The Binarized-Scoring Rule accounts for deviations from risk neutrality and expected utility maximization. For the belief question, we use the opponent frame: *"What is the [respective] probability with which the participant of the preceding experiment you were randomly*

matched to chose the individual boxes of the current set-up?"[11] At the end of the experiment, we randomly select two periods for payment. In one period, the outcome of the game is paid and in the other period, the belief task is paid.

Treatment

Instead of being matched within a session, we matched participants to decisions of earlier sessions. We use data of 360 participants of another study (Bauer & Wolff, 2017) that used the same series of discoordination-games on the same labels. In each period and for every participant of the current study, one decision was sampled from the respective choice distribution (shown in Table C1) of the old experiment. This decision was the payoff-relevant action of "the other player" in the corresponding game. Before playing the game and reporting a belief, participants entered an information stage in which they received varying numbers of observations from the choice distribution they were playing against. The within-subject treatment was n, the number of observations that we sampled and displayed to the participants. The amount of information ranged from 0 to 360 with four periods of zero information.[12] The order of the different n was randomized for $n < 360$ across participants and we informed them that the decision of "the other player" was not contained in the displayed information. In the last period, $n = 360$ for all participants and thus, the information contained the relevant decision (which participants knew).

Manipulation check

After all decisions, we asked participants to indicate their subjective certainty about their belief in three different periods. We showed them the information and their belief report of the periods with $n = \{9, 120, 354\}$. For every of these periods we asked *"How certain are you that your assessment is a good representation of the behavior of your matching partner?"* The certainty was indicated with a slider that ranged from 0 ("absolutely uncertain") to 100 ("absolutely certain") and was not incentivized.

[11] In Bauer & Wolff (2017), we explore the effects of different frames. The opponent frame we use here results in higher belief-action-consistency rates than, for example, a population frame (asking for all other participants' choices). The population frame tends to favour belief-colouring by social projection.

[12] The full set of information levels n was $\{0, 9, 12, 15, 18, 36, 64, 92, 120, 148, 176, 204, 232, 260, 288, 316, 345, 348, 351, 354, 360\}$.

Linking the four-option game to the theory

In the experiment, we opt for a four-option game instead of the two-option game used in our model. We do this for experimental reasons, acknowledging that there are drawbacks when linking our experiment to the model. We prefer the four-option-experiment, because in a two-option game, a randomly clicking person would produce a best-response rate of 50%. Hence observed consistency will be generally high in this case, which makes it likely we would face ceiling effects. In the four-option game, the random best-response rate is reduced to 25%. Also, with four options, we can create much more variance in the label patterns than with two options. Thus, we can keep up participants' interest for more rounds. Further, we also think that participants are more involved in the experiment if they have more influence on their outcome. However, in a two-option game with symmetric payoffs, the influence of a participant's decision on her payoff would be literally minimized.

As indicated above, our model does not trivially extend to the multi-option case. There are more special cases for the prediction of the error rate after updating. In Appendix 3.A we spotlight in an example that the two most important intuitions of our model above carry over to a multi-option case. First, a higher variance of the belief distribution will increase the error rate and larger sample-sizes will *ceteris paribus* decrease the variance. Second, the predictions for Case I i) and Case III i) will hold also in the multi-option case. A more extreme sample will decrease the error rate and a sample with a completely different best-response (and small weight) will increase the error rate. In the simulation described in Appendix 3.B, we implement our model in the four-option-game environment we use in our experiment. The results show that our predictions also hold there.

Procedures

The experiment was programmed using z-tree (Fischbacher, 2007). We use data of 55 participants recruited with ORSEE (Greiner, 2015). All sessions took place in the LakeLab at the University of Konstanz and lasted for approximately 75 minutes, including a short questionnaire at the end of the session which paid 5€. The last item of the questionnaire was a *reliability-of-answers* measure which gives participants the opportunity to indicate how reliable their data is in their opinion. The average payment was 13.27€. See Appendix 3.D for the instructions.

3.4 Predictions

In this section we specify the hypotheses for our experiment. We base the hypotheses on our model predictions in section 3.2.3 and on PROPOSITION 1 which states that the *observed* best-response rate decreases in the error rate ε_k. Given we do not observe directly all variables that are relevant in our theory, we have to proxy for some of them. In the following paragraphs, we discuss the relevant proxies before we turn to the experimental hypotheses.

Approximating congruence of prior and sample

Our theory predictions and hypotheses mostly rely on the relationship of the sample mean μ_s to i) the prior mean μ_q and ii) the prior's relative weight w. Both of these variables are not observed in our experiment. First, we cannot observe the particular strength of participants' prior beliefs $(\alpha + \beta)$, so we do not know w. Second, as the core idea of our theory, participants are not able to access and report μ_q or even $q(\phi)$. For our data analysis we have to rely on proxies for the unobservables.

We proxy the relative weight $w = \frac{n}{\alpha+\beta+n}$ by our treatment variable n, the number of provided observations. This proxy works well for weak priors and loses accuracy in the strength of the prior $(\alpha + \beta)$. It hence could be that a participant by chance gets a high number of observations whenever her prior is particularly strong, and a low number of observations when she has only a weak prior. However, we randomize the treatment n across participants and games. Therefore, there is no reason to expect that such cases will systematically occur or dominate our data.

Second, we proxy situations of a 'congruent sample' and a 'sample in between' by the relationship of the sample to the *reported* belief ϕ^r_B, instead of the relationship of the sample to the prior mean μ_q. We compare what the best-response to both entities separately would be. Hence, we compare on which of the four options the participant places the smallest probability mass in her reported belief to where the minimum number of observations is in the sample.

If the reported belief, ϕ^r_B, has a different minimum than the information and hence also a different best-response, it is highly likely that the information favored a different response than the prior mean, μ_q (Case III), and was not enough to 'overturn' the (reduced) prior. In particular, if participants were able to obtain their true posterior μ_p by some form of sensible updating (including Bayesian updating) and report μ_p, it would *have to be* that the sample contradicted the participant's prior.

In contrast to that, if the reported belief has the same minimum as the information (that is,

$BR[\phi_B^r] = BR[\mu_s]$) it is unlikely that the information differed completely from the reduced prior (Case I & II), unless the information 'overturned' the prior. We will further discuss the influence of 'overturned beliefs' on consistency below when we present our hypotheses.
As a summary, we proxy the relative weight of the information by its number of observations n. The relationship between prior- and sample-mean is approximated by a dummy which compares the *reported* belief to the information. *Belief-min = Info-min* indicates Cases I, II and III ii). Both proxies should work well on average.

Hypotheses

In situations where the sample-information favors the same action as the mean prior belief (Cases I: Congruent sample and II: Sample in between), the expected observed best-response rate always increases in the relative weight. Using our proxies, we can formulate

HYPOTHESIS 1: In cases where participants report beliefs such that *Belief-min = Info-min*, the observed best-response rate increases in the sample size n.

Our situation proxy cannot perfectly separate all cases in which the sample information does not favor the same action as the mean prior belief [Case III]. In particular, Case III ii) consists of cases in which the prior belief is 'overturned' by the information. These cases will also fall into the category *Belief-min = Info-min*. However, beliefs that have just been 'overturned' will have a high variance, which would speak against our Hypothesis 1. We nevertheless expect Hypothesis 1 to hold because we expect these cases to be rare enough not to dominate the data. In any case, not separating these cases from Cases I and II goes against our Hypothesis, so that we should have even more confidence in the effect in case we find it.
Case III i) is indicated by *Belief-min* $\neq$ *Info-min*. Whenever participants report a belief with *Belief-min* $\neq$ *Info-min*, it is highly likely that the belief was not overturned by the sample. This indicates a strong prior. However, because the provided sample differs from the prior, the sample shifts the posterior towards the critical threshold. Hence, in these cases belief uncertainty is generally higher, compared to cases with *Belief-min = Info-min*.

HYPOTHESIS 2A: In cases where participants report beliefs such that *Belief-min* $\neq$ *Info-min*, the observed best-response rate is lower on average, compared to situations with *Belief-min = Info-min*.

2B: If *Belief-min* $\neq$ *Info-min*, the observed best-response rate decreases in the sample size n.

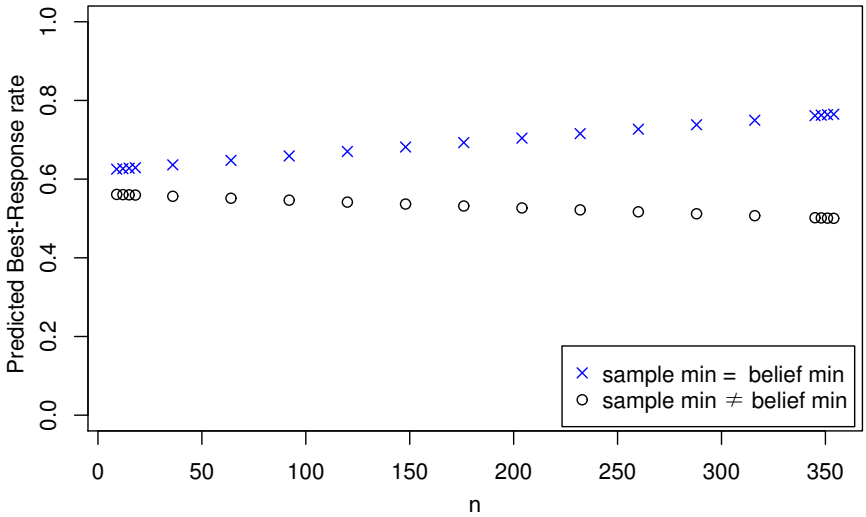

Figure 2: Predicted best-response rates in the four-option game according to our simulation.

For Hypothesis 2B, consider cases with *Belief-min* $\neq$ *Info-min* and relatively large sample sizes. In these instances, even the large n was not sufficient to overturn the prior. We hypothesize that in these cases the belief uncertainty must be particularly high, because posteriors will be close to the critical threshold. Hypotheses 1, 2A and 2B bear out when we simulate the predictions of our theory for the four-option game, as depicted in Figure 2. We describe the setup of the simulation in detail in Appendix 3.B.

3.5 Results

Our most important results are depicted in the left panel of Figure 3, where we use all observations where the belief report has a unique best-response. For each value of our treatment variable n we compute the observed best-response rates across all participants, separately for both values of our situation proxy. If prior and information are not clearly incoherent, the best-response rates are increasing in n (Hypothesis 1) and higher on average compared to situations with contradictory information (Hypothesis 2A). Additionally, in the case when the information clearly contradicts the prior, the best-response rate decreases in n (Hypothesis 2B). These results are statistically supported by a linear regression and Spearman's rank correlations, reported in the right panel of Figure 3.

The results are in line with the predictions of our model. We interpret the different situations created by the interaction of our *Belief-min = Info-min* dummy and our treatment variable n as different levels of belief uncertainty. As predicted in our model, observed best-response rates decrease in belief uncertainty. In the following, we present regressions that also account for

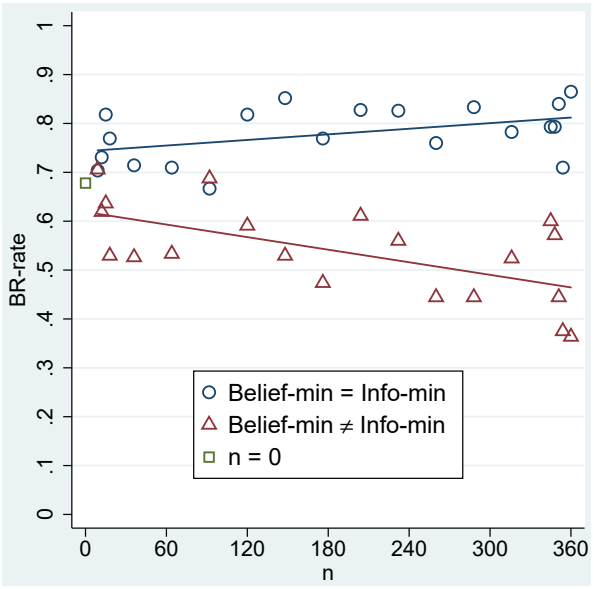

Linear Regression

Best-response rate	Coef.	p-value
$n_{normalized}$	-0.154	0.000
Belief-min = Info-min	0.124	0.001
$n_{norm.}\times$(*Belief-min=Info-min*)	0.223	0.000
Constant	0.619	0.000
	n=40, $R^2 = 0.809$	

Spearman's rank correlations

r_s(Best-response rate, n)	r_s	p-value
Belief-min = info-min	0.460	0.041
Belief-min ≠ info-min	-0.661	0.002

Figure 3: Best-response rates for each n, separated by the relationship of belief and information. $n_{normalized} = \frac{n}{360}$ and $n > 0$ for the regression analysis and Spearman's rank order correlation. There are on average about 27 participants contained in each dot with *Belief-min = Info-min* (blue circles) and 18 per dot in the other group (red triangles).

decision-specific incentives as a robustness check.

Accounting for error costs and learning

The results in Figure 3 use aggregate best-response rates across all participants and hence ignore individual characteristics and incentives. Using regressions that also account for decision-specific incentives, we control for two additional influences on observed best-responses. First, we account for feedback-free learning over time by controlling for the period in which the decision has been made. Second, we account for the cost of making an error. In Section 3.2, we already pointed to the potential effects of Fechner-type decision errors on the observed best-response rate. We account for both factors in the logit regressions whose average marginal effects we report in Table 1.[13] Model 1 tests a model that only includes Fechner-type errors, while Model 2 only includes belief-uncertainty and no error cost (like in Figure 3). Model 3 tests for both sources of errors jointly. Again, we use all observations with a unique best-response and $n > 0$.

Model 1 regresses individual best-responses on individual characteristics and the 'strength' of the belief report ϕ_B^r. By the strength of the belief we mean the utility cost of a decision error as specified by a model with Fechner-type errors, assuming an expected-utility function (which

[13]The results are virtually the same when using the linear-probability model reported in Table C2 in Appendix 3.C.

Best-response to belief	Average marginal effects		
Observations = 898, Clusters = 54	Model 1	Model 2	Model 3
$n_{normalized}$		-0.126**	-0.128**
		(0.051)	(0.050)
Belief-min = Info-min		0.108*	0.108*
		(0.057)	(0.055)
$n_{normalized}$ × *(Belief-min = Info-min)*		0.214**	0.219**
		(0.103)	(0.103)
'Strength' of the reported belief	0.561*		0.576*
	(0.309)		(0.299)
Period	0.009***	0.008***	0.008***
	(0.002)	(0.002)	(0.002)
Male	0.173**	0.167*	0.132*
	(0.076)	(0.076)	(0.073)
Mean Squared Error (Full Sample)	0.1961	0.1869	0.1826
Mean Squared Error (Out of Sample for even Periods)	0.1960	0.1853	0.1814

Table 1: Average marginal effects of logit regressions on observed best-responses. Standard errors in parentheses are clustered on the participant level (54 clusters). The interaction is computed using the `inteff` software by Norton, Wang & Ai (2004). See also Ai & Norton (2003). The marginal effect of the interaction is positive for all participants. Asterisks: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Additional controls in all models: *age*, *math-grade*, *economics-student* and a self reported *reliability-of-answers* measure.

makes the utility cost linear in the probability of a decision error). The strength of a reported belief is thus the percentage-point difference in beliefs on the options with the minimum and the second-lowest probability mass. If the strength is very low, the participant is almost indifferent between choosing the optimal or the second-best option and, according to a model with Fechner-type errors, has a high probability of making such an error. If Fechner-type errors apply, consistency will be low in these situations *independent* of belief uncertainty. The results of Model 1 show that the utility cost of making an error indeed have a large impact on belief-action consistency. High costs of an error strongly increase the probability of an observed best response.

Model 2 replicates our earlier results with respect to belief-uncertainty which hence also hold when accounting for decision-specific incentives. The probability of a best-response decreases in belief uncertainty. Including both sources of error (the strength of belief and belief uncertainty) in the regression shows that the effect of belief-uncertainty is robust also when controlling for the utility cost of an error (Model 3). Finally, feedback-free learning over time leads to more best-responses in later periods in all three models.

To compare all three stochastic-choice specifications, we use out-of-sample predictions. We

perform the regressions in Table 1 for all odd periods and predict the probability of a best response for each decision in all even periods.[14] The bottom panel of Table 1 shows that the out-of-sample mean squared prediction error decreases from Model 1 to 3. To test the predictive power of the models, I compute the average squared prediction error of each model for every subject individually. The distributions of mean prediction errors differ between Model 1 and 2 (Wilcoxon signed-rank test, $p = 0.043$). This means Model 2 outperforms Model 1. Further, Model 3 outperforms both Models 1 and 2 (Model 1 vs 3: $p = 0.009$, Model 2 vs 3: $p = 0.083$). Our results provide evidence that classical decision errors alone cannot explain stochastic choice and belief-action consistency sufficiently in our data. Models 2 and 3, where we add our measures for belief uncertainty clearly outperform the 'standard' decision-error Model 1 both in terms of fit to the data and predictive power. Hence, belief uncertainty plays an important role on top of classical decision errors.

Response times as an alternative measure of utility differences

Above, we use the strength of the reported belief as a measure of the utility cost of an error—hence as a measure for the strength of participants preferences. An alternative measure for the strength of preferences are response times. There is ample evidence in the literature that response times are closely linked to preferences: longer response times indicate that a person is close(r) to indifference between two options.[15] In this study, the response time also may serve as an implicit measure of the strength of preference. This measure might be even less noisy than the strength of the reported belief because it does not rely on the participant's belief report, which, after all, is stochastic according to our model.

We hence rerun our regressions, accounting also for response times. The regressions are reported in Table C3 in the Appendix. We include the normal logarithm of the response time (needed to select and confirm one of the boxes) as an additional explanatory variable in the set of logit regressions reported in Table 1. As expected, the extended models show that quicker response times are associated with higher belief-action consistency. This effect is in line with our above interpretation, that stronger preferences lead participants to committing fewer errors, which in turn leads to higher belief-action consistency. The effect of response times on consistency is robust to adding the belief strength, our original measure of the utility cost of

[14] The out-of-sample results are robust to predicting the choices of the second half of periods (13-24) by the the first half of periods (1-12). However, models 1 and 2 do not differ significantly in that case (Wilcoxon signed-rank test, $p = 0.312$)

[15] Mosteller & Nogee, 1951; Moffatt, 2005; Chabris *et al.*, 2009; Alós-Ferrer *et al.*, 2012; Dickhaut *et al.*, 2013; Konovalov & Krajbich, 2017. Alós-Ferrer *et al.*, 2016 even include this fact as a building block in their economic model to explain preference reversals.

making an error. Most importantly, though, the effect of belief uncertainty is robust to adding response times as an alternative measure for utility differences. Higher belief uncertainty still leads to less belief-action consistency when we include both measures for sources of stochastic choice—belief strength and response times—either separately or jointly. The effect of belief uncertainty becomes stronger, if at all.

The last period with full information

In the last period, where we provided all 360 observations out of the choice distribution, there should be no more (belief) uncertainty about the relevant choice distribution. However, we still do not observe 100% best-responses. Also, 11 participants reported a belief with *Belief-min* $\neq$ *Info-min*. We attribute these observations mainly to other sources of stochastic choice than belief uncertainty. For example, the cost of an error are still relevant when there is no belief uncertainty. Further, it is also conceivable that some participants did not understand that there was no more uncertainty in this period. The data with $n = 360$ are just in line with the rest of the results, as if there was some uncertainty left. All our main results hold (especially all regression results), when excluding the last period with $n = 360$ from the analysis.

Expected earnings in the discoordination game

So far, we have analyzed the effect of belief uncertainty on best-response rates. However, the question remains whether belief uncertainty will cost participants money. If belief uncertainty causes errors in the decisions, participants should earn less if uncertainty is high. We test this hypothesis by looking at the effect of belief uncertainty on expected success rates in the discoordination game. On average, our participants discoordinated in 77.7% of the cases. To assess the performance of our participants in the game, we compute the expected probability to discoordinate given the participant's choice and the true (*i.e.*, complete) choice distribution. In Table 2, we regress the expected probability of discoordinating on our measures for belief uncertainty and the controls.

Model 1 shows that if the participants' prior was congruent with the information, the probability of discoordinating increases in the number of observations in the sample, as expected. The effect is practically nil if prior and information were not congruent. But what should we expect given our model? *A priori*, the answer to this question is unclear.

When prior and information are incongruent, there will be three partially counterveiling effects on our observed variables. First, note that providing information which contradicts prior beliefs will do two things: it will increase choice stochasticity—which is good if you would otherwise

Expected probability of discoordinating	Model 1	Model 2
$n_{normalized}$	0.001 (0.005)	0.002 (0.005)
Belief-min = Info-min	0.012*** (0.003)	0.011*** (0.003)
$n_{normalized}$ × *(Belief-min = Info-min)*	0.023*** (0.007)	0.021*** (0.006)
Relative distance of belief to true distribution		-0.026*** (0.009)
Period	0.0004** (0.0002)	0.0004** (0.0002)
Male	0.002 (0.004)	0.003 (0.003)
Constant	0.749*** (0.028)	0.757*** (0.029)

Table 2: Linear regressions of the expected probability to discoordinate, given the true choice distribution. Standard errors in parentheses are clustered on the participant level (54 clusters). Asterisks:*** $p < 0.01$ ** $p < 0.05$. Additional controls: *age*, *math-grade*, *economics-student* and a self reported *reliability-of-answers* measure. The expected probability to discoordinate ranges from 65.8% to 83.1% in the data.

always choose the wrong option—and it will make the posterior more adequate than the prior. Both effects would mean performance should increase in the amount of provided information n also for *Belief-min* $\neq$ *Info-min* observations. However, there also will be a selection effect. For low n, there will be both people with high and people with low relative weight on the prior in the *Belief-min* $\neq$ *Info-min* group. Among this group, the people with high relative weight on the prior will perform worse, because they nearly always choose the wrong option. In contrast, people with low relative weight will sometimes choose the right option because of the variance in their belief distribution. If we now increase the amount of information, people with low relative weight will tend to drop out of the *Belief-min* $\neq$ *Info-min* observations. In other words, for increasing n, the better-performing people will drop out of the average, which is the selection effect counterveiling the two performance-increasing effects of increasing n.[16] What the data seems to show is that the counterveiling effects seem to just cancel out on average.

In Model 2, we want to estimate how much of the positive effect of more information is due to the decrease in belief uncertainty, and how much is due to more accurate beliefs (*i.e.*, to μ_p being closer to the true ϕ^*). To do so, we additionally control for how close the participants

[16]The people dropping out of the *Belief-min* $\neq$ *Info-min* average will enter the *Belief-min = Info-min* average, of course. There, they will bring down the average because they will be the least likely to choose the right option in this group. In other words, our estimation will underestimate the beneficial effect of more information for both groups.

n	Mean normalized certainty	Std. Dev.	Rank-sum test
9	-5.45	16.62	$p = 0.020$
120	0.10	12.35	$p = 0.016$
354	5.35	16.90	

Table 3: Results of the unincentivized certainty-question in three different rounds. The certainty report is normalized by individual means.

belief report was to the true choice distribution. If participants receive more observations from the true distribution, their belief is shifted more and more towards the true distribution and hence becomes more accurate, the larger *n* is. A more accurate belief however, should increase the probability to discoordinate independently of belief uncertainty. To control for participants' more accurate beliefs, we include the distance of their belief report to the true distribution as a control variable.

Model 2 shows that both the improved accuracy in beliefs and the reduction of belief uncertainty play an important role in improving expected payoffs. The effect sizes are roughly the same for changing from the maximum possible difference between belief and true distribution to reporting the true distribution and for changing from providing virtually no information to all potential information, at least for the *Belief-min* = *Info-min* case. Note again, though, that the coefficients for both variables containing $n_{normalized}$ will be biased downwards due to the selection effect described above.

In summary, we find that lower belief uncertainty is associated with a higher probability to discoordinate provided that prior beliefs are not inaccurate. In turn, high belief uncertainty causes participants to forgo actual money in the experiment in that case. At the same time, belief uncertainty quite naturally will be beneficial when beliefs are inaccurate, as it will move participants away from invariably choosing the wrong thing.

Unincentivized certainty

Table 3 shows the result of the unincentivized certainty questions at the end of the experiment. For each subject, the three certainty reports for their beliefs in rounds with $n \in \{9, 120, 354\}$ are normalized by the participant's mean certainty level, to level out individual heterogeneity. On average, the reported certainty increases in the amount of information and the distributions differ significantly across n according to rank-sum tests. These results further support our interpretation of uncertain beliefs and that on average our manipulation of certainty was meaningful.

3.6 Conclusion and Discussion

In many cases, people play according to equilibrium only after sufficient experience.[17] In this paper, we point out that experience with a situation may matter not only for whether people play equilibrium strategies—it matters also for whether they act optimally given their (unobserved) belief distribution. If people find themselves in a completely new situation, it is very likely that they 'don't really know what to believe' about the uncertain features of the situation. When people face this type of environment, we propose that their choices will exhibit a large variance that depends not only on their average belief and the expected costs of making mistakes like in other models, but on the degree of belief uncertainty people face.

We model belief uncertainty by the variance in players' probability distributions over possible beliefs. This belief uncertainty creates stochastic choice because players have no direct access to their belief distribution, so that players have to sample a belief each time they need to act. Thereby, our model and experimental evidence point to a source of stochastic choice that so far has been neglected in the literature. Taking belief uncertainty into account will be important when predicting people's choices in situations where they face high degrees of strategic or environmental uncertainty. However, it remains open to further research how belief uncertainty plays out in more sophisticated decision problems: for example, when we buy a house, we hesitate in order to think about it multiple times. In our model, this makes sense if we revisit the decision time and again to sample more probabilities. An interesting question that ensues here is, of course, how we integrate those sampled probabilities. Does multiple sampling decrease belief uncertainty? To test this conjecture, the experimenter would have to make participants think about a relevant probability several times and only then require a choice and a belief report.

As a final exercise in this paper, we show that belief uncertainty relates to expected earnings. Also in this regard, our model provides a new perspective. When beliefs are at least somewhat accurate, increasing belief uncertainty will cost people money. While this is not predicted by standard theory, it is hardly surprising. What is less obvious is that—always controlling for the reduced belief—belief uncertainty can be beneficial, namely when beliefs are inaccurate. This may be a reason for why people might tend to entertain a relatively high degree of uncertainty about their beliefs: when it is not clear whether my belief is accurate or not, a high degree of uncertainty acts as a hedging device—at least I will do the right thing some of the time.

Belief uncertainty is also highly relevant for us as researchers when eliciting beliefs and in-

[17] *E.g.*, Fudenberg & Levine (2016), and references cited therein.

terpreting experiment participants' belief-action consistency. Consistency will be predictably low when participants face a new, unknown, and complex situation; it will be higher, the more participants get to know the situation. In particular, belief-action consistency will be higher the more participants learn what others will do. This again links to the literature on whether experiment participants learn to play the Nash equilibrium. Research has shown that the degree of complexity is an important factor (*e.g.*, Grimm & Mengel, 2012). Our study highlights why more complex environments are often associated with less equilibrium play. Not only may the beliefs be far from being equilibrium beliefs, but the associated belief uncertainty will often make the other equilibrium concept's key ingredient fail: the ingredient of best-response behavior.

References

Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80(1), 123-129.

Alós-Ferrer, C., Granić, D. G., Shi, F., & Wagner, A. K. (2012). Choices and preferences: Evidence from implicit choices and response times. *Journal of Experimental Social Psychology*, 48(6), 1336-1342.

Alós-Ferrer, C., Granić, D. G., Kern, J., & Wagner, A. K. (2016). Preference reversals: Time and again. *Journal of Risk and Uncertainty*, 52(1), 65-97.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593-1636.

Bauer, D. & Wolff, I. (2017). Biases in Beliefs: Experimental Evidence. *Mimeo*.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1963). Stochastic models of choice behavior. *Behavioral Science*, 8(1), 41-55.

Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity Analysis of Production and Allocation*, 13(1), 374-376.

Camerer, C., & Hua Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4), 827-874.

Chabris, C. F., Morris, C. L., Taubinsky, D., Laibson, D., & Schuldt, J. P. (2009). The Allocation of Time in Decision-Making. *Journal of the European Economic Association*, 7(2/3), 628-637.

Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies*, 75(3), 729-762.

Cubitt, R. P., Starmer, C. & Sugden, R. (2001). Discovered preferences and the experimental evidence of violations of expected utility theory. *Journal of Economic Methodology*, 8(3), 385-414.

Dickhaut, J., Smith, V., Xin, B., & Rustichini, A. (2013). Human economic choice as costly information processing. *Journal of Economic Behavior & Organization*, 94, 206-221.

Dow, J., & da Costa Werlang, S. R. (1992). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60(1), 197-204.

Etner, J., Jeleva, M., & Tallon, J. M. (2012). Decision theory under ambiguity. *Journal of Economic Surveys*, 26(2), 234-270.

Fechner, G. (1860/ 1966). Elements of Psychophysics, Vol. 1, New York: Holt, Rinehart and Winston

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

Fudenberg, D., & Levine, D. K. (2016). Whither game theory? Towards a theory of learning in games. *The Journal of Economic Perspectives*, 30(4), 151-169.

Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 141-153.

Gilboa, I., Postlewaite, A. W., & Schmeidler, D. (2008). Probability and uncertainty in economic modeling. *The Journal of Economic Perspectives*, 22(3), 173-188.

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

Grimm, V., & Mengel, F. (2012). An experiment on learning in a multiple games environment. *Journal of Economic Theory*, 147(6), 2220-2259.

Groeneveld, R. A., & Meeden, G. (1977). The mode, median, and mean inequality. *The American Statistician*, 31(3), 120-121.

Guiso, L., & Parigi, G. (1999). Investment and demand uncertainty. *The Quarterly Journal of Economics*, 114(1), 185-227.

Harless, D. W., & Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica*, 62(6), 1251-1289.

Hey, J. D., & Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 62(6), 1291-1326.

Hossain, T., & Okui, R. (2013). The binarized scoring rule. *The Review of Economic Studies*, 80(3), 984-1001.

Konovalov, A. & Krajbich, I. (2017). "Revealed Indifference: Using Response Times to Infer Preferences." Working Paper.

Loomes, G., & Sugden, R. (1995). Incorporating a stochastic element into decision theories. *European Economic Review*, 39(3), 641-648.

Manski, C. F., & Neri, C. (2013) First- and second-order subjective expectations in strategic decision-making: Experimental evidence. *Games and Economic Behavior*, 81, 232-254.

McKelvey, R. D., & Palfrey, T. R. (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10(1), 6-38.

Moffatt, P. G. (2005). Stochastic choice and the allocation of cognitive effort. *Experimental Economics*, 8(4), 369-388.

Mosteller, F., & Nogee, P. (1951). An experimental measurement of utility. *Journal of Political Economy*, 59(5), 371-404.

Norton, E. C., Wang, H., & Ai, C. (2004). Computing interaction effects and standard errors in logit and probit models. *Stata Journal*, 4, 154-167.

Plott, C.R. (1996) 'Rational individual behavior in markets and social choice processes: the discovered preference hypothesis', In K.J. Arrow, E. Colombatto, M. Perlman and C. Schmidt (eds) *The Rational Foundations of Economic Behaviour*, Basingstoke: Macmillan, pp. 225–50.

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscienc*e, 19(3), 366-374.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59.

Savage, L. J. (1954) *The Foundations of Statistics*. New York: John Wiley and Sons. (Second ed., Dover, 1972).

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3), 571-587.

Selten, R., & Stoecker, R. (1986). End behavior in sequences of finite Prisoner's Dilemma supergames A learning theory approach. *Journal of Economic Behavior & Organization*, 7(1), 47-70.

Weber, Roberto A. (2003). 'Learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, 44(1): 134-144.

Appendix

3.A Intuition of our model with three actions

Assume a game with three options $X = \{X_1, X_2, X_3\}$. For the best response to a belief, it is only relevant which of the options the agent believes to be the least likely choice by the other player. Every possible belief on such a game is a vector of three probabilities $\phi = \{\phi_1, \phi_2, \phi_3\}$ (one for each option) with $\sum_{j=1}^{3} \phi_j = 1$. That means every possible belief is a point on the 2-simplex. A belief distribution (over all possible beliefs) is hence a multivariate probability distribution $v(\phi)$ with the 2-simplex as support.[18] This belief distribution assigns a (subjective) probability to every possible "three-way" belief (point on the 2-simplex). Also with the multivariate belief distribution, the player faces a compound lottery. She might for example, assign a probability of 40% to $\phi'' = \{0.55, 0.1, 0.35\}$. Standard theory would however assume a best response to the reduced lottery and hence to the *expected* choice probabilities $E_v[\phi] = \{E[\phi_1], E[\phi_2], E[\phi_3]\}$.

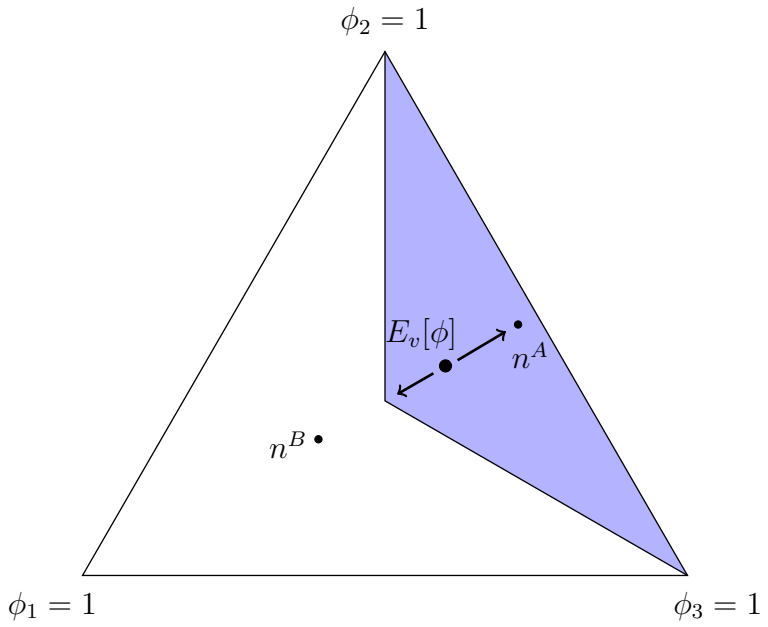

Figure A1: 2-Simplex for all possible beliefs in a three-option game. Exemplary expected choice probabilities and two observed samples n^A and n^B

Suppose the best response to $E_v[\phi]$ was x_1. This means, that $E[\phi_1]$ is the smallest element of $E_v[\phi]$. This situation is depicted in Figure A1. The shaded area indicates all beliefs with best response x_1.

Assume the agent reacts to one draw ϕ^r from $v(\phi)$ instead of best responding to the expected

[18] A natural assumption is the Dirichlet-distribution, the multivariate generalization of the Beta-distribution. This distribution is a conjugate prior for the multinomial distribution.

probabilities. Then she commits an error, whenever she draws a ϕ^r that has a different minimum element than $E_v[\phi]$. The error rate is hence characterized by the probability mass on all possible beliefs that have a different minimum than $E_v[\phi]$. Such beliefs are depicted by the non-shaded area in the simplex.

In our two-option model above, the error rate increased in the variance of the belief distribution. This also holds for the multi-option case. Suppose the distribution $v(\phi)$ is spread out around $E_v[\phi]$ with high variance. Then it is more likely, that the player draws a ϕ^r with a different minimum than $E_v[\phi]$. Now, how is the error rate further influenced, for example by updating? To see this, denote $E_v[\phi] = \mu_v$ as the "prior mean" and suppose the agent observes one of two possible samples n^A, n^B. The maximum-likelihood of a specific sample can be depicted as one point on the simplex. The intuition is the same as in the two option case: After bayesian updating, the new posterior mean $E_z[\phi] = \mu_z$ will be a weighted sum of the sample and the prior mean. This is also indicated by the arrows in Figure A1. The (posterior) mean of the belief distribution will be pulled towards the sample points on the simplex. Now consider the sample n^A. It is more extreme, compared to $E_v[\phi]$ (Like Case I in Section 3.2.3). The posterior mean is pulled towards the edge of the simplex and more probability mass of the belief distribution is shifted to the shaded area. This means that the error rate will decrease.

The opposite happens when the agent observes sample n^B (Similar to Case III). It favors a different best response and the posterior mean is pulled towards the center of the simplex. This means that probability mass of the belief distribution is pushed on the non-shaded area and the error rate increases. All these patterns are further moderated by the posteriors variance which, as in the two-option case, decreases in sample size.

3.B Simulating the four-option game

To make clear that our predictions for the four-option game do indeed result from our theory, we run a simulation. First, we randomly choose an absolute weight n_q for our prior, with $n_q \sim U[1, 400]$. n_q can be interpreted as the number of observations in a prior sample. We choose an upper limit of 400 so that there can be priors that outweigh the maximum sample size of 360 used in our experiment. Our prior should be Dirichlet-($\boldsymbol{\alpha}$) distributed (*cf.* ftn. 18). So, we randomly draw four probabilities $\pi_i^{(q)}$ for the αs of the prior distribution. We use a Dirichlet-$(1, 1, 1, 1)$ distritution for this random draw. Then, we use the randomly-drawn probabilities together with the drawn n_q, to determine the parameters of the prior Dirichlet distribution: $\alpha_i^{(q)} = n_q \pi_i^{(q)} + 1$.

After randomly defining the prior, we create an "observed sample" of choices. We draw the number of new observations n from a uniform distribution over all levels we use in the experiment but the extreme cases, so that $n \sim U\{9, 12, 15, 18, 36, 64, 92, 120, 148, 176, 204, 232, 260, 288, 316, 345, 348, 351, 354\}$. Then, we randomly determine 'choice probabilities' for the random samples. For this purpose, we draw three values $\pi_i^{(aux)}, \pi_i^{(aux)} \sim U[0, 1]$. We then let sampling probabilities be a random perturbation of the following sequence of probabilities: $\pi_1^{(s)} = \pi_1^{(aux)}$, $\pi_2^{(s)} = (1 - \pi_1^{(s)})\pi_2^{(aux)}$, $\pi_3^{(s)} = (1 - \pi_1^{(s)} - \pi_2^{(s)})\pi_3^{(aux)}$, and $\pi_4^{(s)} = (1 - \pi_1^{(s)} - \pi_2^{(s)} - \pi_3^{(s)})$. Using the random perturbation of our probabilities $\pi_i^{(s)}$, we draw a sample of n new 'observations'. We then apply Bayesian updating to update the prior Dirichlet distribution according to the 'new observations', so that $\alpha_i^{(p)} = \alpha_i^{(q)} + n_i$.

So far, we have simulated a prior belief-distribution with absolute weight n_q and an observed sample of n choices. Using Bayes' rule, we have updated the prior to arrive at a posterior belief distribution. To assess the predicted observed best-response rate for the resulting posterior, we use 10'000 iterations of the following process: from the posterior, we draw a belief ϕ_A^r for the action and a belief ϕ_B^r for the reported belief, with $\phi_A^r, \phi_B^r \sim Dir(\boldsymbol{\alpha^{(p)}})$. If the two beliefs have their minimum on the same option, they are consistent. For each draw of ϕ_B^r, we also record whether it has the same minimum as the distribution of 'new observations' $\boldsymbol{n}$. Then, we record the average consistency for all draws of ϕ_B^r that have the minimum on the 'anti-mode' of $\boldsymbol{n}$. Further, we compute the average consistency for all draws of ϕ_B^r that do not have the minimum on the 'anti-mode' of $\boldsymbol{n}$ (where we define the anti-mode to be the location that occurs least often in the sample). We thus compute best-response rates separately for when the reported belief indicates the same best-response as the observed sample and when it has not.

We iterate the above process 5'000 times. Then, we use a linear regression to relate the level of consistency to the sample size n, a dummy indicating whether the drawn belief ϕ_B^r has its minimum on the anti-mode of the sample $\boldsymbol{n}$, and the interaction of both terms. We plot the resulting predicted best-response rates in Figure 2 in Section 3.4. This prediction has three characteristics: when the reported belief and the sample suggest the same choice, (i) the best-response rate is higher than when they do not; (ii) the predicted best-response rate increases in n; and (iii) when the reported belief and the sample suggest different choices, the predicted best-response rate decreases in n.

3.C Figures and Tables

Game	Box 1	Box 2	Box 3	Box 4	χ^2	Sig. on 5%	Sig. on 1%
1	74	106	110	70	14.578	✓	✓
2	110	68	76	106	14.844	✓	✓
3	84	70	86	120	15.022	✓	✓
4	110	100	70	80	11.111	✓	-
5	104	84	101	71	7.933	✓	-
6	76	77	97	110	9.044	✓	-
7	115	63	84	98	16.156	✓	✓
8	83	90	87	100	1.7556	-	-
9	123	74	75	88	17.489	✓	✓
10	104	83	92	81	3.667	-	-
11	97	77	81	105	5.822	-	-
12	101	82	88	89	2.111	-	-
13	86	76	80	118	12.178	✓	✓
14	116	92	72	80	12.267	✓	✓
15	76	104	89	91	4.378	-	-
16	91	66	102	101	9.356	✓	-
17	113	70	90	87	10.422	✓	-
18	85	95	61	119	19.244	✓	✓
19	100	76	71	113	13.178	✓	✓
20	93	87	75	105	5.200	-	-
21	97	84	86	93	1.222	-	-
22	92	71	93	104	6.333	-	-
23	102	75	101	82	6.156	-	-
24	104	67	76	113	16.111	✓	✓
Number of significantly non-uniform distributions:						15	10

Table C1: The 24 historic choice distributions, used to sample the provided information. Corresponding χ^2-tests with H_0: choices are uniform across boxes

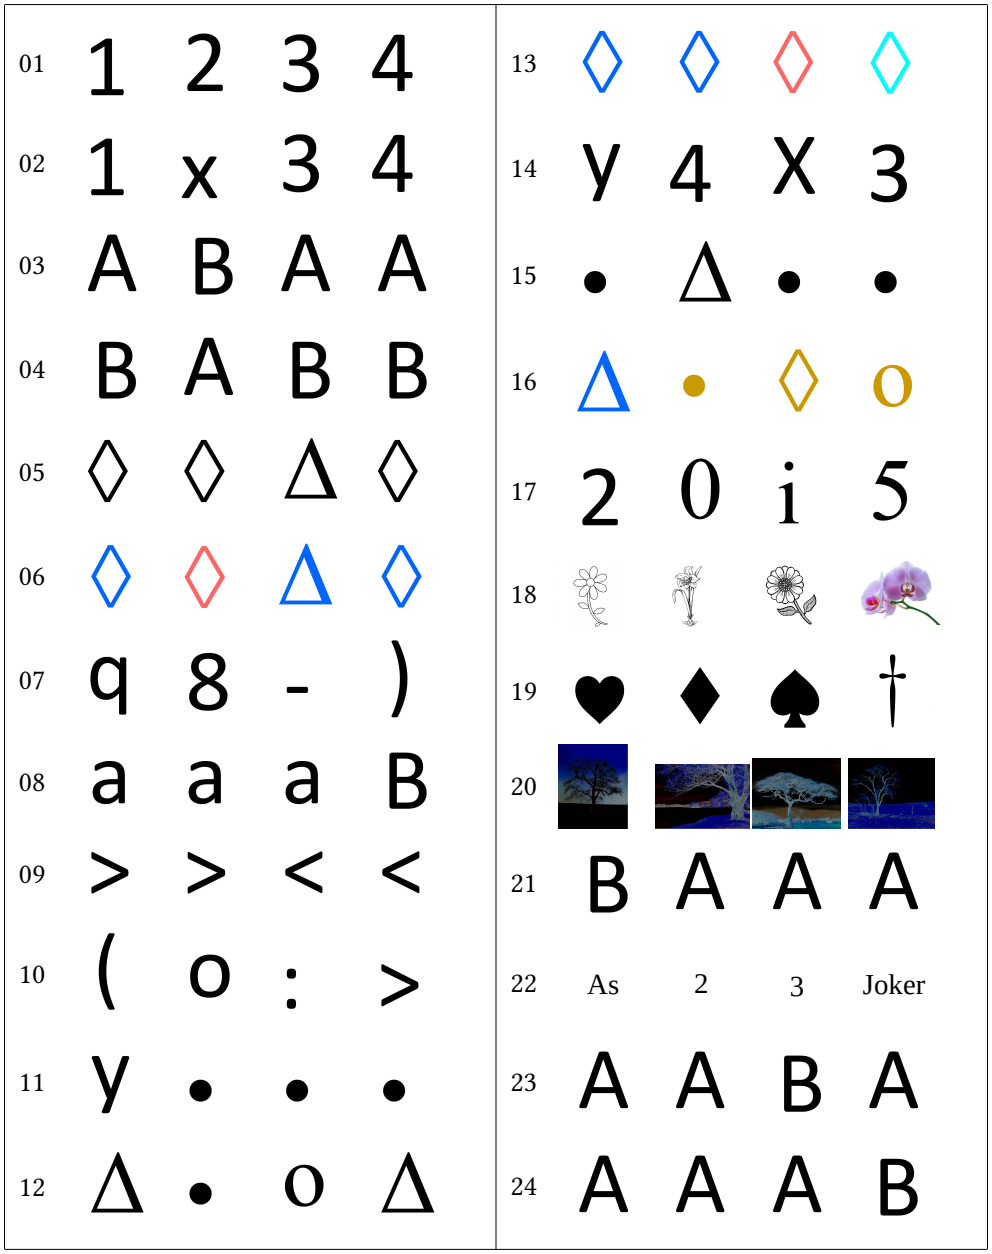

Figure C1: The 24 label sets, used to label the four options of the games. One set for each game. Sources of the pictures from sets 18 and 20 can be found in the picture credits.

Best-response to belief	Linear Probability Model		
	Model 1	Model 2	Model 3
$n_{normalized}$		-0.158**	-0.160**
		(0.068)	(0.066)
Belief-min = Info-min		0.117*	0.117*
		(0.063)	(0.061)
$n_{normalized}$ × *(Belief-min = Info-min)*		0.212**	0.217**
		(0.102)	(0.100)
'Strength' of the reported belief	0.511**		0.534**
	(0.223)		(0.222)
Period	0.009***	0.008***	0.008***
	(0.002)	(0.002)	(0.002)
Male	0.169**	0.157**	0.128*
	(0.069)	(0.069)	(0.066)

Table C2: Linear Probability Model OLS regressions of observed best-responses. Standard errors in parentheses are clustered on the participant level (54 clusters). Asterisks: *** p<0.01, ** p<0.05, * p<0.1. Additional controls in all models: *age*, *math-grade*, *economics-student* and a self-reported *reliability-of-answers* measure.

Best Response to belief	Average Marginal effects after Logit			
	Model 1′	Model 1″	Model 2′	Model 3′
ln(decision time)	-0.123***	-0.114***	-0.099***	-0.090***
	(0.033)	(0.034)	(0.033)	(0.034)
$n_{normalized}$			-0.124**	-0.128**
			(0.051)	(0.050)
Belief-min = Info-min			0.095*	0.094*
			(0.054)	(0.053)
$n_{normalized}$×*(Belief-min = Info-min)*			0.211**	0.217**
			(0.103)	(0.102)
'Strength' of the reported belief		0.490		0.516*
		(0.300)		(0.299)
Period	0.006***	0.007***	0.006***	0.006***
	(0.002)	(0.002)	(0.002)	(0.002)
Male	0.181**	0.155**	0.149**	0.121*
	(0.077)	(0.075)	(0.073)	(0.072)
Mean Squared Error	0.1944	0.1914	0.1834	0.1799

Table C3: Marginal effects of Logit regressions accounting for ln(decision time). Number of Observations = 898. Standard errors in parentheses are clustered on the participant level (54 clusters). Asterisks: *** p<0.01, ** p<0.05, * p<0.1. Additional controls in all models: *age*, *math-grade*, *economics-student* and a self-reported *reliability-of-answers* measure.

3.D Experimental Instructions

The instructions are translated from german. Boxes indicate consecutive screens showed to participants.

Today's Experiment
Today's experiment consists of 24 rounds in which you will make two decisions each.

Decision 1 and Decision 2

In the first round, you will see the instructions for both decisions directly before the decision. In later rounds, you can display the instructions again if you need to.

The payment of the experiment
In every decision you can earn points. At the end of the experiment, 2 rounds are randomly drawn and payed. In one of the rounds, we pay the point you earned from decision 1 and in the other round, you earn the points from decision 2. The total amount of points you earned will be converted to EURO with the following exchange rate:

1 Point = 1 Euro

After the experiment is completed, there will be a short questionnaire. For completion of the questionnaire, you additionally receive 5 Euro. You will receive your payment at the end of the experiment in cash and privacy. No other participant will know how much money you earned.

General Instructions
For todays experiment, another experiment plays a central role. This experiment has been conducted earlier, here in the LakeLab. The earlier experiment is describet in the following.

The earlier experiment

In the earlier experiment, 360 participants ran through 24 rounds. In every round groups of two randomly matched persons were formed. The group members did not know each others identity and could not communicate throughout the whole experiment.
One round of the experiment worked in the following way: both participants did see the exact same screen. On the screen, there was an arrangement of four boxes which are marked with symbols. Both of the group members chose one of the boxes. If both group members chose different boxes, both received a price. If both members chose the same box, there was no payoff. All participants learn about which box was chosen by the other participant and which payoff they received in a certain round only at the end of the experiment.
The arrangement of symbols on the boxes differed in every round for every group. The decision of a participant was hence on an unknown arrangement. Below, you can see an example of how such an arrangement could have looked like.

Example: The four boxes are marked from left to right by Diamond, Heart, Spade, Diamond.

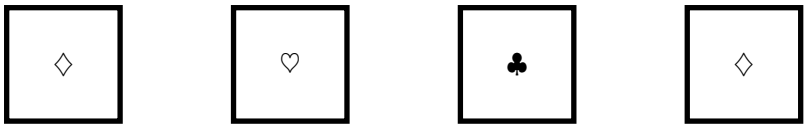

In this example, there are two boxes which are marked with the same symbol. However, the boxes on the most left and most right count as are different boxes.

Instructions for experiment 1
The number of points you receive in decision 1 depends on your own decision, as well as on a participant of the earlier experiment who will be randomly matched with you. How this works, will be explained in the following.

Decision 1

For decision 1 in every round, you see an arrangement of four boxes which are marked with symbols that was also used in the earlier experiment. The computer then randomly draws one of the participants of the earlier experiment who chose one of the boxes.

In decision 1, you have to choose a box as well.

If you choose another box than your randomly matched partner from the earlier experiment, you receive 7 points. If you choose the same box as your randomly matched partner, you don't receive points.

On the next screen, you receive more information about the earlier experiment.

Additional Information
In every round, before you make decision 1, you receive additional information, how a certain sample of the 360 participants of the earlier experiment decided in the respective arrangement. In every round, a random sample is drawn from all 360 participants of the earlier experiment. For every of the four boxes, you get to know how many participants in the sample chose that box. You can see an example of how this information looks like below:

[Example Screen, see screenshot below]

Please note, that the participant you are matched to in the respective period is not contained in the sample you see. This means, that this participant is always drawn from the remaining participants which are not shown to you.
The size of the respective sample of participants you receive information about will vary from round to round. This means, that you have different amounts of information about the decisions of the participants of the earlier experiment in every round.
Please note, that the participants of the earlier experiment did not have any information how other participants decided. Information like you can see it above, was not displayed to the participants of the earlier experiment.
The information is displayed on the next screen.

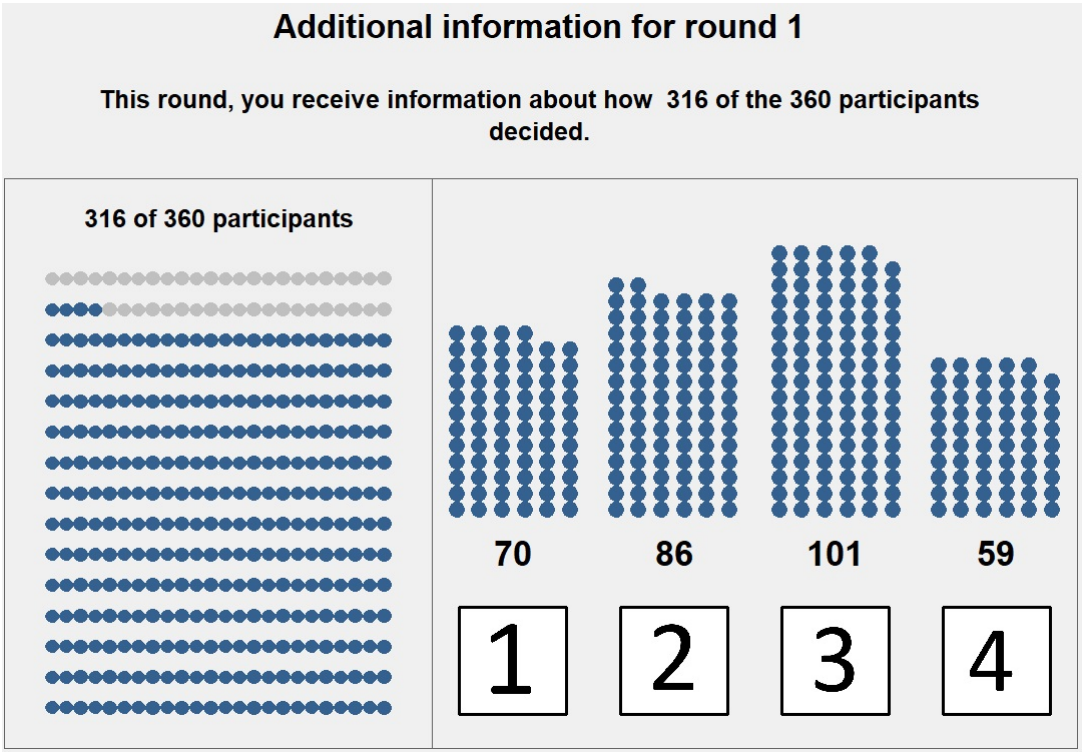

Instructions for decision 2
In decision 2, your payoff also depends on your own decision and on the decision of your matching partner from the earlier experiment. We now explain decision 2 in detail.

Decision 2
Decision 2 refers always to the arrangement from decision 1, which was also used in the earlier experiment. You will hence see the arrangement of boxes from the respective round again. You also can look at the additional information again. Again, the decision of your matching partner from the earlier experiment is relevant for you.
Decision 2 is about your assessment, how your matching partner from the earlier experiment decided. We are interests in your assessment of the following question:

"With what probability did your matching partner chose each of the respective boxes of the current set-up?"

For every box, you can report your assessment with what probability your matching partner chose the respective box. You can enter the percentage numbers in a bar diagram. By clicking into the diagram, you can adjust the height of the bars. You can adjust as many times as you like, until you confirm.
Since your assessments are percentage numbers, the bars have to add up to 100%. The sum of your assessment is displayed on the right. You can adjust this value to 100% by clicking. Or you enter the relative sizes of your assessments only roughly and then press the "scale" button. Please note, that because of rounding, the displayed sum ma deviate from 100% in some cases.
On the next page, we explain the payoff of decision 2.

The payoff in decision 2
In this decision, you can either earn 0 or 7 points. Your chance of earning 7 points increases with the precision of your assessment. Your assessment is more precise, the more it is in line with the decision behavior of your matching partner. For example, if you reported a high assessment on the actually selected box, your chance increases. If your assessment on the selected box was low, your chance decreases.
You may now look at a detailed explanation of the computation of your payment, which rewards the precision of your assessment.

It is important for you to know, that the chance of receiving a high payoff is maximal in expectation, if you assess the behavior of your matching partner correctly. It is our intention, that you have an incentive to think carefully about the behavior of your matching partner. We want, that you are rewarded if you have assessed the behavior well and made a respective report.

At the end of the experiment, one participant of today's experiment will roll a number between 1 and 100 with dies. If the rolled number is smaller or equal to your chance, you receive 7 points. If the number is larger than your chance, you receive 0 points.
As soon as you reported and confirmed your assessment about the behavior of your matching partner, the round ends. You will then be matched with another participants and the next round begins.

Payment of the assessments
At the end of your assessment, you will receive the 7 points with a certain chance (p) and with $(1-p)$, you receive 3 points. You can influence your chance p with your assessment in the following way:

As described above, you will report an assessment for each box, on how likely your matching partner is to select that box. One of boxes is the actually selected. At the end, your assessments are compared to the actual decision of your matching partner. Your deviation is computed in percent.

Your chance p is initially set to 1 (hence 100%). However, there will be deductions, if your assessments are wrong. The deductions in percent are first squared and then divided by two.

For example, if you place 50% on a specific box, but [your matching partner selects another box,] your deviation is equal to 50%. Hence, we deduct $0.50{*}0.50{*}\frac{1}{2} = 0.125$ (12.5%) from p.

[For the box, which is actually selected by your matching partner, it is bad if your assessment is far away from 100%. Again, your deviation from that is squared, halved and deducted. For example if you only place 60% probability on the actually selected box, we will deduct $0.40 * 0.40 * \frac{1}{2} = 0.08$ (8%) from p.]

With this procedure, we compute your deviations and deductions for all boxes.
At the end, all deductions are summed up and the smaller the sum of squared deviations is, the better was your assessment. For those who are interested, we show the mathematical formula according to which we compute the chance.

$$p = 1 - \frac{1}{2}\left[\sum_i (q_{box_i,estimate} - q_{box_i,true})^2\right]$$

The value of p of your assessment will be computed and displayed to you at the end of the experiment. The higher p is, the better your assessment was and the higher your chance to receive 7 points (instead of 0) in this part. At the end of the experiment, the computer will roll a random number between 0 and 100 with dies. If this number is smaller or equal to p, you receive 7 points. If the number is larger than p you receive 0 points.

Summary
In order to have a high chance to receive the large payment, it is your aim to achieve as few deductions from p as possible. This works best, if you have an good assessment of the behavior of your matching partner and report that assessment truthfully.

Chapter 4

Responder Communication in Ultimatum Bargaining[§]

Dominik Bauer & Fabian Dvořák

Graduate School of Decision Sciences & Thrurgau Institute of Economics

University of Konstanz, Germany

Abstract: We theoretically and experimentally study responder communication in the ultimatum game. Before the game starts, responders send a structured and non-binding message about the smallest offer, they are willing to accept. Our theoretical results show, that information transmission through communication is limited in equilibrium. In line with this result, we find that responders largely send the message which suggests to split the pie equally, even though they would also accept smaller shares of the pie in the large majority of cases. Our elicited beliefs show, that i) despite the result before, responders send lower messages than their beliefs would indicate, ii) communication is not cheap talk, because proposers' beliefs are correlated with messages and iii) also proposers deviate from what their beliefs suggest to do. They do not increase offers further for very high messages. The observed discrepancy of beliefs and actions for both player roles is in line with the predictions of inequality aversion, but only if we acknowledge the existence of extreme preference types.

JEL classification: C78; C91; D84; D91

Keywords: Communication; Ultimatum game; Reciprocal behavior; Social preferences

4.1 Introduction

Empirical observations show, that there exists considerable heterogeneity in behavior in the ultimatum game (see Roth, 1995 and Güth & Kocher, 2014 for reviews). While some people even accept the smallest possible offer, fairly large shares of the pie are rejected by others. Social preference models explain this variation in behavior by assuming different parameters

[§]Financial support of the Graduate School of Decision Sciences is gratefully acknowledged. Contact: dominik.3.bauer@uni-konstanz.de, fabian.dvorak@uni-konstanz.de. Graduate School of Decision Sciences, University of Konstanz, Box 146, D-78457 Konstanz, Germany.

in the utility function of these individuals. In this paper, we present a variant of the ultimatum game in which the responder can send a cost-less, non-binding message before play. In the message, responders can convey what the minimum offer is they are barely willing to accept. We theoretically and experimentally study the extend to which responders communicate their preferences over bargaining outcomes to the proposer.

Theory suggests that efficiency should increase if responders' minimum acceptance thresholds were public information. The better the proposer is informed about the preferences of the responder, the more likely will she make an offer which will be accepted. Independent of the players' preference types, there always exists an allocation of payoffs which generates more utility for each player compared to rejection. However, our theoretical analysis shows that truthful revelation of the acceptance thresholds is difficult to achieve, because the increase in efficiency comes at a cost for the responders, who typically receive less if they reveal their preferences. As a consequence there exists a upper bound on the amount of information responders can transmit in equilibrium.

Apart from this bound on information transmission, a theoretical analysis of the game suffers the common problem of multiple equilibria in games with incomplete information. If we focus purely on players actions, nearly every outcome of the game can be theoretically justified by fixing (arbitrary) beliefs. This makes it difficult to judge the validity of theory based on observations of actions only.

Thus, we complement our theoretical result of limited information transmission with an empirical investigation. We conduct an experiment with a standard ultimatum game and implement structured communication to the game in two treatments, in which responders can either send a message about what they are minimally willing to accept or what they perceive as an "justified offer". Apart from the communication, a main feature of our experiment is the belief elicitation. Throughout the whole experiment, we make extensive use of the strategy method (Selten, 1967), for both actions and beliefs. For example, we elicit the proposer's posterior beliefs about the minimum acceptance threshold of the responder, conditional on the message. We do not restrict the elicitation of beliefs to the messages which are actually sent, but also elicit beliefs for messages the theorist would consider off the equilibrium path. Our participants hence report beliefs for every message that could potentially be sent. Such belief data can help us to empirically tackle the problem of multiple equilibria and relate the observed behavior to the predictions of inequality aversion.

The experiment unveils that there is indeed only limited information transmission. There is substantial pooling on messages around 5, but these messages are often inflated relative to re-

sponders rejections thresholds. However, messages are not uninformative. Proposers' beliefs about responders bargaining preferences differ over the message space. Messages which indicate high responder demand are believed to come from responders with higher acceptance thresholds. However, in contrast to the prediction of standard theory, offers do not increase along with the proposers' beliefs in the upper half of the message space. For high messages, proposers offer less than what their standard selfish best-response (corrected for risk preferences) suggests. This discrepancy of proposers' actions and beliefs is in line with the model by Fehr & Schmidt (1999), once we extend the type space usually assumed by the model. We find evidence that proposers believe in the existence of very demanding bargaining types, which do not exists according to the standard parameter assumptions of the model.

Related Literature

Other experimental papers have studied the effects of cheap-talk communication in the ultimatum game. While we study responder communication, Andersson *et al.* (2010) study the effect of proposer messages and find that these allow proposers to obtain a higher payoff from the game. Boles, Croson & Murnighan (2000) study two-way communication over private information and document that both proposers and responders are willing to use deceptive strategies (see also Croson, Boles & Murnighan, 2003 and Lundquist *et al.*, 2009). Zultan (2012) studies the difference between game-related and non-game-related open-chat communication and finds that only game-related communication changes subjects' strategies. The most closely related paper is Rankin (2003) who also employs one-way responder communication. He finds a positive correlation between non-biding responder requests and offers. At the same time, the requests decrease the average payoff of the game by reducing offers on average, compared to a no-communication treatment.[1] The focus of Rankin (2003) is to identify the effect of communication on the observable outcomes of the game. However, without beliefs it is difficult to judge why the messages have certain effects. Further, the "requests" carry some normative element in our opinion. In the light of the common social preferences models, it is unclear what effect such an normatively loaded message has. Our paper hence contributes to the literature in two ways. First, we focus on explaining *why* the messages affect the outcomes by eliciting participants on- and off-equilibrium path beliefs. Second, in our main communication treatment, the message has no direct normative meaning, but carries only the intention to act in a certain way. In an additional treatment, we also study messages with the normative meaning of what offer is considered as *justified*.

[1] We thank David Rankin for kindly sharing his data with us.

4.2 Predictions and research questions

Results of the ultimatum game show that people differ in their preferences over bargaining outcomes. We assume these inter-individual differences can be accounted for by a single number, which is ex-ante private information: the smallest amount a responder is willing to accept. In our setup with communication, the responder can potentially transmit information about her minimum acceptance threshold to the proposer. This information can affect the proposer's belief about the responder's behavior.[2]

We use the model of Fehr & Schmidt (1999) for a theoretical analysis of the standard ultimatum game with structured communication. We introduce communication to the model by assuming that the responder can send a message about her minimum acceptance threshold to the proposer. If there exists heterogeneity in the model's parameters, the message can potentially contain information about the responder's social preference type, for example her (dis-) advantageous inequality aversion (Fehr-Schmidt's α_i & β_i). We assume that the distribution of social preference types in the population is common knowledge but that each individual only privately observes her type. The main result of the theoretical analysis is captured in Proposition A, which suggests that there exists a limit of information transmitted in equilibrium:

> Proposition A: In the Fehr-Schmidt model, there is no sequential equilibrium in the ultimatum game with communication, in which responders fully separate according to their type if
> (1) the distribution of disadvantageous-inequality-aversion parameters α_i is such that there exist at least two different rejection thresholds and
> (2) there exists a proposer type which is only weakly averse to advantageous inequality with $\beta_i < \frac{1}{2}$.
> Proof: Appendix 4.B

The intuition is simple. Because there is heterogeneity in social preferences, some types will accept relatively low offers as responders (e.g. an offer of 2) but the same offer would be rejected by other types. As an example, we pick two specific types from the distribution and call the type accepting [rejecting] low offers the *low* [*high*] type. Now assume a low-type responder revealed herself to a proposer by transmitting information about her type via communication.

[2] If one assumes no variation in social preference parameters, communication should not have an effect since no game-related private information can be exchanged. For the same reason communication should not have an effect according the standard model. Every responder accepts any positive amount and hence every proposer offers the smallest positive amount.

Under condition (2) of Proposition A, there exist proposer types which will offer the lowest amount the revealed responder is willing to accept. Now further assume that a high-type responder reveals himself to the same proposer (type). Also the high type will - with positive probability - receive the lowest amount he is willing to accept. However, this lowest amount will be higher than the one which was offered to the low type. The low type has hence an incentive to mimic the high type and to pretend to reject low offers. Hence, there is no equilibrium in which responders fully separate according to their type.

Following from Proposition A, responders must send messages in equilibrium, which yield the highest possible payoff. Hence, different types use the messages in the same way, *i.e.* the distribution of used messages is independent of the type. For our experiment derive:

> Hypothesis 1: Responders' messages are uncorrelated with the acceptance thresholds.

Apart from Hypothesis 1, there are many questions which cannot be answered based on theory alone, mostly because there are multiple equilibria.[3] We conduct the experiment to tackle the following questions: (1) What messages are used by responders? (2) To what extend do responders transmit information about their preferences through communication? (3) How are messages interpreted by proposers? How do proposers' on and off-equilibrium beliefs look like? Further, if information is transmitted, (4) to what extend is behavior affected by the information, for example through social preferences?

To shed light on the hypothesis and research questions, our experiment makes use of the strategy method. In order to approximate the distribution of responder types in our sample, responders decide for each possible offer, whether to accept or to reject by indicating a minimum acceptance threshold. We use the strategy method also for belief elicitation in order to characterize off-equilibrium beliefs. For example, we also elicit beliefs about rejection thresholds after extreme messages which should rarely (or never) be actually sent. By eliciting out-off equilibrium beliefs, we can investigate whether the observed behavior is in line with the predictions of inequality aversion.

[3] There are many equilibria where no information is transmitted on the equilibrium path. It is quite intuitive that any equilibrium in the game without communication has many equivalents in the game with communication, in which every message received in equilibrium is uncorrelated with the responders' type. One example is the babbling equilibrium where every responder type sends every possible message with the same probability. In all of these equilibria proposers' beliefs (and hence also their behavior) are unaffected by the communication. On the other hand, a large number of outcomes can be supported as equilibria by fixing beliefs in response to out-of-equilibrium messages. For instance, each distribution of used messages can be rationalized by assuming that proposers interpreted all unused messages as being exclusively sent by the least demanding responder type.

The distribution of messages will answer question 1. To answer question 2, we look for a correlation between these thresholds and the message sent. This informs us about the amount of information which is contained in responders' messages. For question 3, we check how much information is transmitted to the proposer. We compare proposers' beliefs about the responder's acceptance threshold for all messages (sent and not sent). Question 4 is certainly the hardest one to answer since we cannot specify all model parameters based on our data. However, we use the qualitative prediction of inequality aversion to investigate systematic deviations from standard best-responses. We calculate subjects standard best-response based on elicited beliefs and risk-preferences and check whether deviations in the observed behavior deviates are in line with the predictions of social preferences.

4.3 Experiment

4.3.1 The ultimatum game

Our experiment uses a standard one-shot, two-player ultimatum game. In this game, one of the players (the proposer) has to make an offer on how to allocate 10 monetary units (MU) between himself and the other player. The other player (the responder) subsequently chooses to accept and implement the allocation or to reject the offer, in which case both players receive zero MU. We introduce communication to the ultimatum game in two different treatments. Before the ultimatum game begins, the responder can send a structured, cost-less and non-binding message to the proposer. In the treatments, we vary the context of the communication by a between-subjects manipulation of the framing of the message. In one treatment, the message is framed as information about what the smallest offer is, the responder is willing to accept (MIN treatment). In another treatment, the message is framed as information on what the responder perceives as *justifiable offer* (JUST treatment). Responders can select any message from the set $\{0, 1, ..., 10\}$. For the actual decisions in the game, we use the strategy method.[4] Without knowing which message has been selected by the responder, proposers have to select an offer from the set $\{0, 1, ..., 10\}$ for every possible message that can be sent. At the same time, without knowing which offer has been selected for their message, responders decide for any offer that can be made whether to accept or to reject by selecting a minimum offer from $\{0, 1, ..., 10\}$ they are willing to accept. In the following we will refer to this as the acceptance threshold of the responder or the minimum. We also run a control treatment without communication to check

[4]We are aware of the fact that the use of the strategy method can influence game play (Oosterbeek, Sloof & Van De Kuilen, 2004). However, eliciting participants' full strategy vector is a crucial feature of our design.

if our treatments influence the outcome of the game.

4.3.2 Beliefs

Before our participants learn about the result of the ultimatum game, we elicit participants' beliefs about the decisions of the other player. We do this, conditional on each message. This means that the proposer indicates - for each possible message - a belief distribution on how likely each possible acceptance threshold is. For example, for a message of 5, a proposer may assign a probability of 40% to a threshold of 5, 30% to a threshold of 4, 30% to a threshold of 3 and 0% to all other thresholds. In total, each proposer reports 11 messages $\times$ 11 threshold-probabilities = 121 subjective probabilities.

In a similar fashion, the responder indicates for each possible message, how likely each possible offer in response to this message is. Although the responder knows which message she actually selected, we ask responders to indicate beliefs for all possible messages, also those they did not send.[5] Each proposer hence reports 11 messages $\times$ 11 offer-probabilities = 121 subjective probabilities.

The belief elicitation was incentivized using the Binarized-Scoring Rule (McKelvey & Page, 1990; Hossain & Okui, 2013). This rule uses a quadratic scoring rule to assign participants lottery tickets for a given prize. The lottery procedure accounts for deviations from risk neutrality and, under a weak monotonicity condition, even for deviations from expected utility maximization.[6] To avoid the possibility of hedging, participants knew that either the ultimatum game or the belief elicitation would be randomly selected for payment.

4.3.3 Additional tasks

We also elicited estimates for participants' risk preferences by a procedure used in Sheremeta & Shields (2013). In this task, participants had to make 15 decisions to choose between a safe payoff of 1 MU and a lottery with varying expected payoff in a list format (see Table A1 in the Appendix). We did not force participants to have a unique switching point within the decisions. One of these decisions was randomly selected for payment at the very end of the experiment. Finally, we elicited socio-economic control variables in a standard questionnaire format. Participants received a lump-sum payment of 2 Euro for completing the questionnaire.

[5] This procedure can be incentivized because we used the strategy method to elicit an offer for each message.

[6] Of course the proper elicitation of beliefs conditional on events which are not believed to occur has its limitations. For example, eliciting a belief distribution based on a message which is not believed to be sent at all, cannot be incentivized properly.

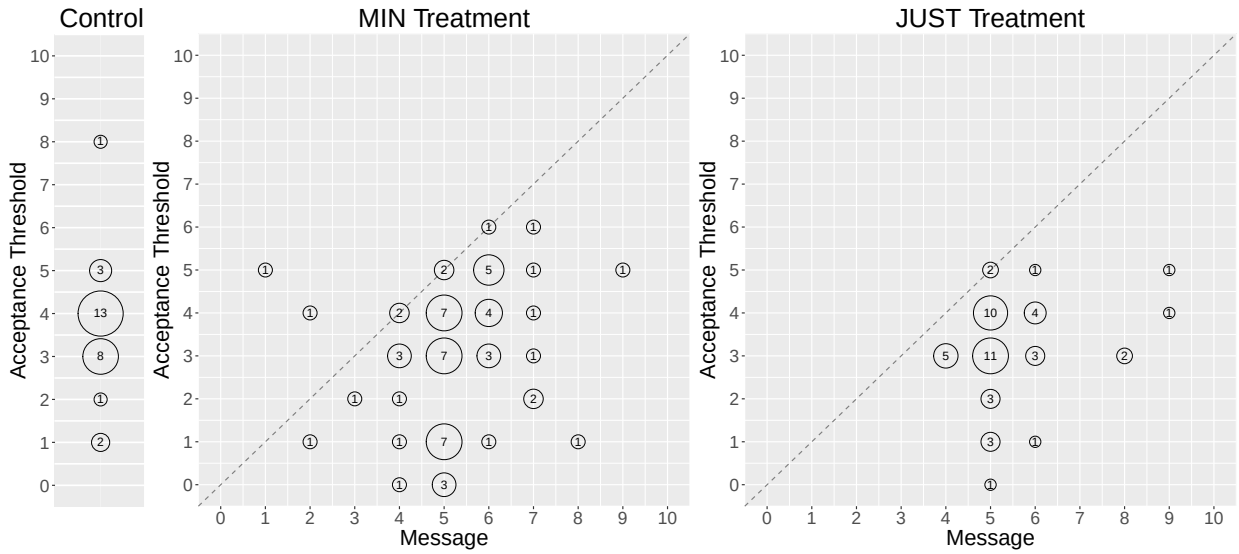

Figure 1: Messages and acceptance thresholds of responders. Bubble sizes and numbers indicate the observed frequency of the respective message-threshold combination. $N_{Control} = 28$, $N_{MIN} = 60$, $N_{JUST} = 48$

4.3.4 Procedures

The experiment was programmed with z-Tree (Fischbacher, 2007). Participants were recruited using ORSEE (Greiner, 2015). We conducted 9 experimental session (5 MIN,4 JUST, 2 Control) in the LakeLab at the University of Konstanz between December 2014 and August 2016 including a total of 272 subjects (120 MIN, 96 JUST, 56 Control). An experimental sessions lasted approximately one hour and subjects received a payment of 14.63 Euro on average. The exchange rate between MU and Euro was 1 MU = 1.5 Euro in all treatments. The instructions which were handed out to participants are reported in Appendix 4.C.

4.4 Results

4.4.1 Responders

Messages and acceptance thresholds

In the experiment, we observe substantial pooling on messages around 5 (Question 1). Figure 1 shows, that 5 is the modal message in both treatments. Further, the framing of the message significantly influences which message is selected by responders. A message of 5 is sent more often in the JUST treatment (62.5%) compared to the MIN treatment (43.3%, Fisher's exact test: $p = 0.055$). Figure 1 also shows, that communication carries some, but only limited infor-

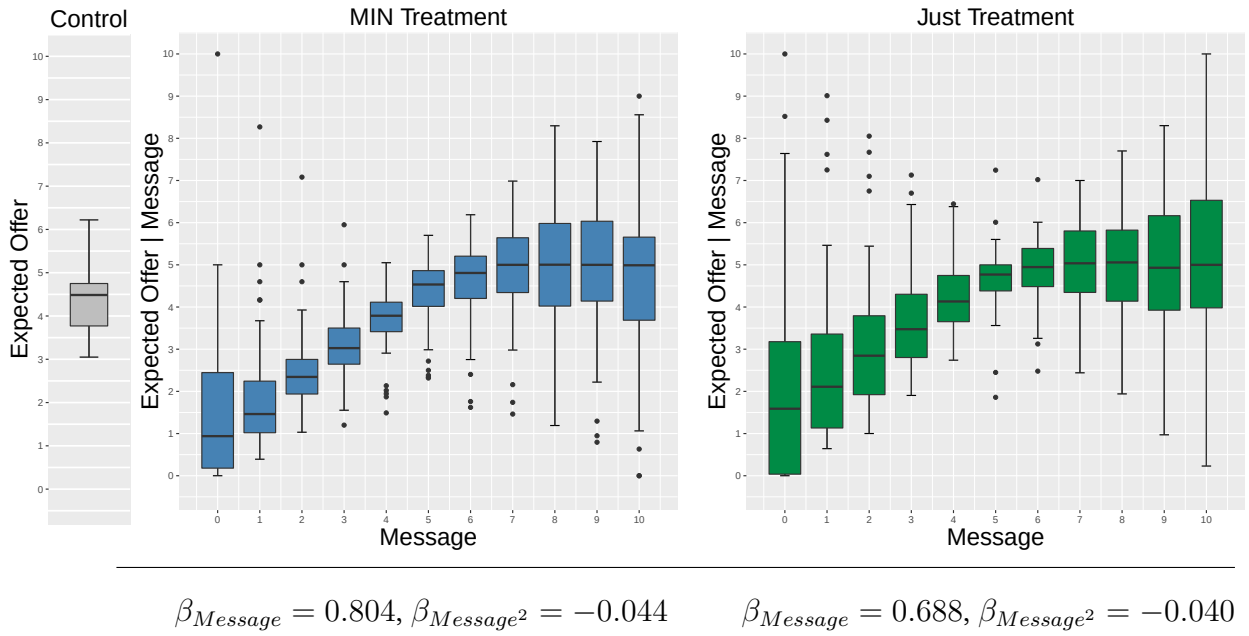

Figure 2: Boxplots of responder beliefs. Coefficients of OLS regressions of expected offer on messages. $p < 0.001$ for all coefficients, standard errors clustered on participant level. $N_{Control} = 28$, $N_{MIN} = 60$, $N_{JUST} = 48$, $R^2_{MIN} = 0.448$, $R^2_{JUST} = 0.278$

mational content with respect to the acceptance threshold. Messages and thresholds are positively correlated (Spearman's rank correlation; MIN: $r_s = 0.298$, $p = 0.021$; JUST: $r_s = 0.242$, $p = 0.090$), rejecting Hypothesis 1. Further, messages are substantially inflated compared to actual thresholds (Question 2). On average, messages are about 2.2 MU larger than actual thresholds in both treatments. In the MIN treatment, about 88% of participants inflated their message. In the JUST treatment, almost 96% of participants inflated their message (Signed-rank tests, $p < 0.001$ in both treatments).[7]

Responder beliefs

Figure 2 shows boxplots of responder beliefs across treatments.[8] Responders expect offers to increase with higher messages, however the positive relation breaks down for high messages. This observation is backed up by clustered OLS regressions of expected offers on messages and squared messages which are also reported in Figure 2. In the upper half of the message space, responders do not believe that higher messages are associated with substantially higher offers.

[7]There is no significant difference in the distributions of messages, thresholds or inflation between the communication treatments. Rank-Sum tests: $p = 0.952$, $p = 0.916$ and $p = 0.937$ respectively. Rejection thresholds in the control treatment (mean: 3.68) are not significantly different compared to the MIN treatment (mean: 3.05, rank sum test: $p = 0.168$) and only differ on a 10 % level to the JUST treatment (3.19, $p = 0.090$).

[8]All boxplots show the median as black line, the upper and lower quartile as colored bars, whiskers indicate the $1.5\times$ interquartile range and dots show outliers.

However, responders do on average expect that, for example a message of 7 yields a higher offer compared to a message of 5.[9] There is hence some discrepancy in responder beliefs and actions. The belief data suggests, that proposer should send high messages. However, the modal observed message is 5. Next we will analyze whether observed actions (*i.e.* sent messages) are best-responses to the reported beliefs.

Best-responses that account for risk preferences

In order to answer question (4) and to make inferences about possible social-preference motivations of our participants, we compute best-responses to beliefs with the assumptions of standard selfish behavior, taking also risk aversion into account. To estimate participants risk preferences, we use the data from the lottery choice task at the end of our experiment. Since we did not require to enter a unique switching point, we estimate the most likely switching point for every participant separately. To this end, we simultaneously estimate a random tremble parameter and the most likely switching point for each participant using maximum-likelihood. For the most likely switching point, we compute the parameter of a Constant-Relative-Risk-Aversion (CRRA) utility function. Using this utility function, we then compute best-responses to reported beliefs. For responders, we determine the message which gives maximal expected CRRA-utility. We call this message the best-response to the responders whole set of beliefs.
We do the same procedure also for proposers. For each message, we compute the offer that gives maximal CRRA-utility, given the expected rejection thresholds. This means, the responder has 11 best-response offers, one for each possible message.
In the following, we will use the term "best-response" for the actions obtained by the above procedure. Of course, this "best-response" does not capture other possible motivations for behavior or deviations from expected-utility maximization. We compare the difference between the standard best-response (accounting for risk aversion) and the actual behavior of participants and relate the difference to the qualitative predictions of the social preferences. Therefore we do not include such motivations in the best-response.

Responder deviation from best-responses.

As the above analysis suggests, Figure A1 in the Appendix shows that responders' messages are lower than their best-responses. On average, participants send a message that is about 2.7 MU lower than their best response (MIN treatment. 1.4 MU in the JUST treatment. Signed-rank tests: $p < 0.003$ in both cases). Our participants deviate strongly from their individual

[9] Means in the MIN treatment: 4.32 vs 4.88, rank sum test: $p < 0.001$. JUST, means: 4.64 vs. 5.05, $p = 0.022$.

standard best-responses and send lower messages. In the MIN treatment, 85% of responders deviate downwards from their best-response. In The JUST treatment 66.7% deviate to a lower message.

4.4.2 Proposers

Figure 3 on page 95 shows box-plots of proposers' offers and beliefs for both treatments and for each message.[10] The top panel shows, that offers increase on average in the message for the lower half of the message space. However, there is also a substantial number of relatively high offers for low messages. Further, the positive relationship breaks down for messages larger than 5. Behavior is hence affected by messages, but not throughout the whole message space. The positive trend, which fades out for high messages is supported by linear regressions of offers on messages and squared messages which are also reported in Figure 3. Similar to the pattern we observed for responder beliefs, proposers do not further increase their offers if the messages are to high (*i.e.* for messages larger than 5) in both treatments. When using only Messages > 5, offers do even slightly *decrease* in the message (Clustered OLS regressions; MIN: $\beta_{Message} = -0.082$, $p = 0.096$; JUST: $\beta_{Message} = -0.085$, $p = 0.201$).

Surprisingly, the belief data in the lower panel of Figure 3 does not fully support the pattern of offers. On average, proposers belief that a higher message is associated with a higher rejection threshold, even for messages larger than 5 (Question 3). The coefficients of OLS regressions of expected thresholds on messages also show a weakening effect for high messages. However, the beliefs still increase, also for messages larger than 5.[11]

If the message is larger than 5, 30% of reported beliefs have an expected threshold of larger than 5 in the MIN treatment. This means, that for high messages some proposers expect the responder to actually reject an offer of 5. However, as already shown above, these beliefs do not translate into behavior. In only about 8% of cases would proposers actually send an offer larger than 5 if the message is larger than 5 (in less than 5% of cases would proposers send an offer of larger than 6).

In the following, we analyze whether also proposers deviate from their best responses. We do so first conditional on messages, then conditional on the size of the best-response and lastly analyze both influences jointly.

[10] The offers in the control treatment (mean: 4.32) are not significantly different compared to the actually sent offers (after the strategy method inputs have been realized) in the MIN treatment (mean: 4.33, rank sum test: $p = 0.630$) and the JUST treatment (4.63, $p = 0.188$).

[11] Clustered OLS regressions of E(threshold) on messages, for messages > 5. MIN: $\beta_{Message} = 0.204$, JUST: $\beta_{Message} = 0.194$. Both $p < 0.003$.

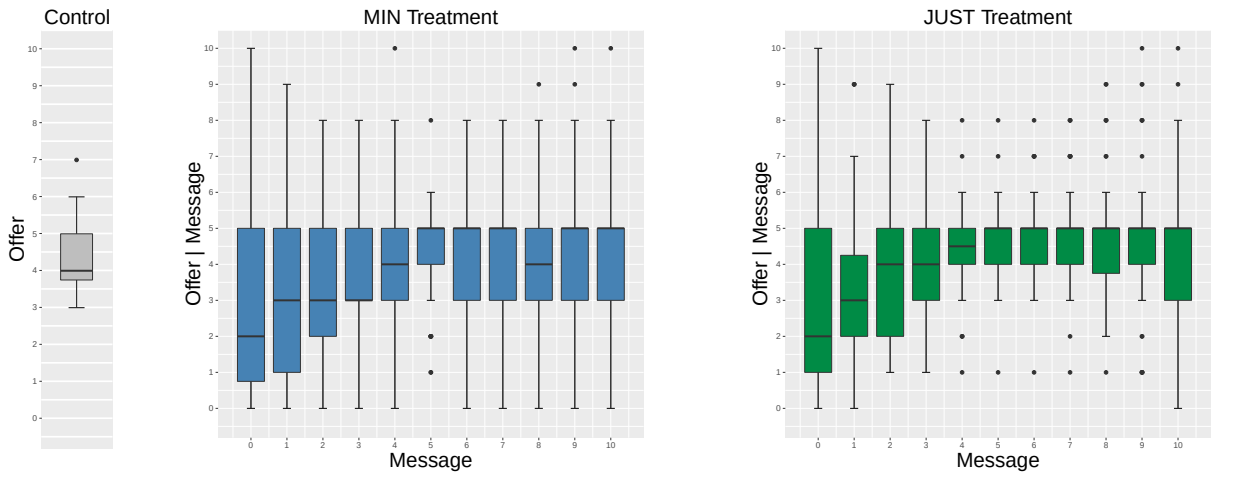

$\beta_{Message} = 0.514, \beta_{Message^2} = -0.038$ $\beta_{Message} = 0.559, \beta_{Message^2} = -0.040$

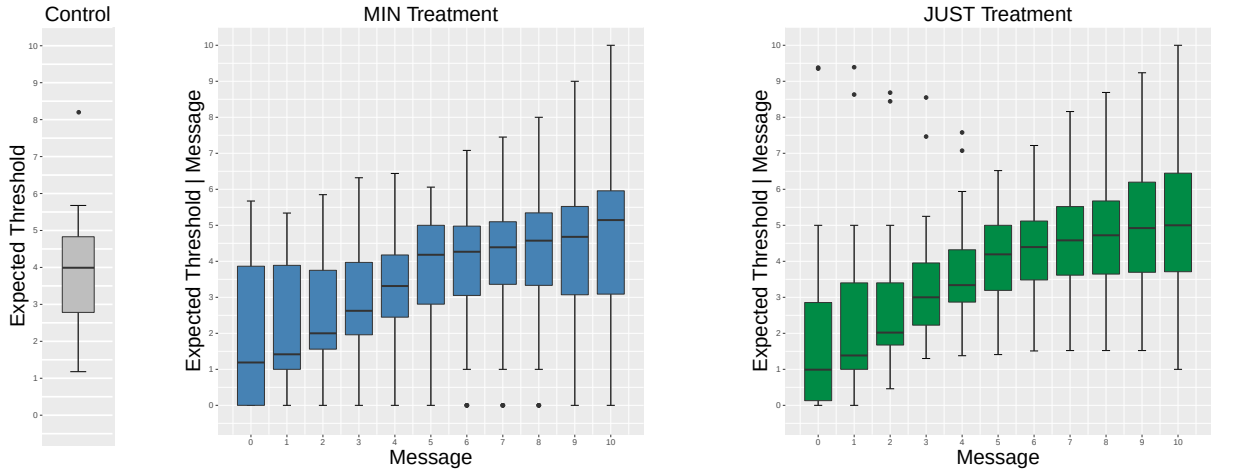

$\beta_{Message} = 0.415, \beta_{Message^2} = -0.014$ $\beta_{Message} = 0.563, \beta_{Message^2} = -0.023$

Figure 3: Boxplots of proposer offers (top panel) and expected rejection thresholds (bottom panel), conditional on the message. Coefficients of OLS regressions of expected offer on messages. $p < 0.004$ for all coefficients with standard errors clustered on participant level.$N_{Control} = 28$, $N_{MIN} = 60$, $N_{JUST} = 48$, Offers: $R^2_{MIN} = 0.079$, $R^2_{JUST} = 0.100$, Beliefs: $R^2_{MIN} = 0.208$, $R^2_{JUST} = 0.291$

Offer - Best-Response	Message										
	0	1	2	3	4	5	6	7	8	9	10
MIN Treatment	0.03	0.20	0.15	-0.02	-0.18	-0.22	-0.37	-0.50	-0.67	-0.60	-1.02
JUST Treatment	0.38	0.13	0.25	0.02	-0.17	-0.48	-0.63	-0.79	-0.88	-1.08	-1.10

Table 1: Mean deviations from best-responses conditional on messages. Page's trend tests: MIN treatment: $L = 25110, p < 0.0001$; JUST treatment: Page's $L = 20182.5, p < 0.0001$

Proposer deviation from best-responses

We compute proposers' best-responses using the same procedure as for responders to asses systematic deviations from those best-responses. For each message, we subtract the individual best-response from the actual offer made. Table 1 shows the average deviation from the best response for each message. The average difference between offer and best-response gradually decreases as the message increases for both treatments and is negative for messages larger than 3. Page's trend tests confirm that the negative trends are statistically significant. This means our participants deviate stronger from their best-response, the higher the message is. These results are backed up by the linear regressions reported in Table A2 in the Appendix, which also account for individual characteristics.

Offers relative to best-responses

Above we show, that deviations from best-responses depend on the message. However, since also the beliefs of proposers substantially vary with the message, it is unclear whether the message has an effect itself, or whether the message serves only as a moderator for beliefs. As proposed by the theoretical model, instead of reacting to the literal meaning of the message, proposers react to the belief about the responders' rejection threshold which is conveyed by the message. In the following, we first analyze how the offers are influenced by best-responses *disregarding* which specific message did lead to resulting beliefs and best-responses. As a last step, we jointly analyze how the two channels (beliefs and messages) affect offers.

Figure 4 shows bubble plots of offers, separately for each best-response.[12] Note that now, the offers are displayed independently of the message. Hence, the number of observations differ across best-responses.[13] Figure 4 shows for both treatments, that if the best response is smaller than 5, most offers are equal or larger than the respective best-response. In this range, 71%

[12]The same data is shown as boxplots in Figure A2 in the Appendix.

[13]By the use of the strategy method, each individual boxplot in Figures 2 and 3 contained the same number of observations. This is not the case in Figure 4. It is for example possible, that a participant reported beliefs such that for all messages, her best-response is equal to 4. All observations of this participant would hence be contained in bubbles with *Best-Response = 4*

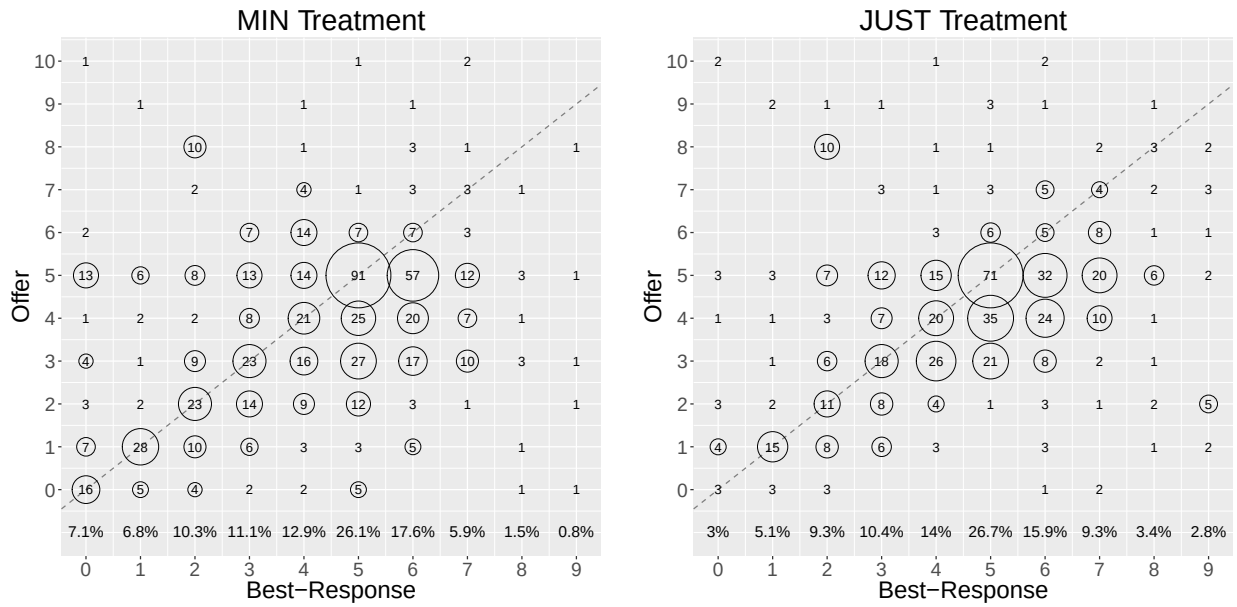

$\beta_{BR} = 1.115^{***}$, $\beta_{BR^2} = -0.090^{***}$ $\beta_{BR} = 0.797^{*}$, $\beta_{BR^2} = -0.050$

Figure 4: Bubble plots of offers conditional on best-responses. Bubble sizes and numbers indicate the observed frequency of the respective Offer - Best-Response combination. Relative frequencies of best-responses are indicated below. Clustered OLS regressions using observations wit best-responses > 0. Additional controls in all models: *male*, *age*, *available money per month*, and a self reported *reliability-of-answers* measure. We further control for the quality of our risk preference estimation by adding the tremble parameter and the standard deviation of the likelihood function. Asterisks: *** p<0.01, * p<0.1. $N_{MIN} = 60$, $N_{JUST} = 48$, $R^2_{MIN} = 0.158$, $R^2_{JUST} = 0.149$

percent of offers in the MIN and 67% of offer in the JUST treatment are larger or equal to the best-responses.

Best-responses larger than 5 are not as frequent, making up only 25% of the data in the MIN, and 32% of the data in the JUST treatment. In these subcases however, only about 12% of the offers are larger than the best-responses in the MIN treatment (about 14% of offers in the JUST treatment). Hence not only do offers vary with the message, but also with the best-response

Regressions of offers on best-responses and squared best-responses in Figure 4 show, that offers increase along with best-responses as one would expect. However, in the MIN treatment, offers *decrease* in the best-response for high best-responses. This means that if beliefs (after receiving a message) are such that the responder has a high rejection threshold, proposers do the opposite to what their best-responses would suggest. In the JUST treatment, the coefficient of the squared term is negative as well, but not significantly so.

Offers	Treatment	
	MIN	JUST
Message	0.246**	0.384**
	(0.115)	(0.172)
Message^2	-0.019**	-0.030**
	(0.008)	(0.013)
Best-Response	1.009***	0.491
	(0.351)	(0.527)
Best-Response^2	-0.083**	-0.026
	(0.035)	(0.049)
Constant	1.595	0.600
	(1.520)	(1.866)
Clusters	58	48
R^2	0.1706	0.1748

Table 2: Linear regressions of offers on messages and best-responses if best-response > 0. Standard errors in parentheses are clustered on the participant level. Asterisks: *** p<0.01, ** p<0.05. Additional controls in all models: *male*, *age*, *available money per month*, and a self reported *reliability-of-answers* measure. We further control for the quality of our risk preference estimation by adding the tremble parameter and the standard deviation of the likelihood function.

The joint influence of messages and beliefs on offers

We now jointly analyze how the two channels of beliefs and messages affect offers. Table 2 shows linear regressions of offers on messages and best-responses, the respective squared terms and controls. In the MIN treatment, both the belief and the message channel influence offers. Proposers do offer less than their best-response suggests for high best-responses, but also react aversely to high messages.

In the JUST treatment, we find no significant evidence for the belief channel when controlling also for messages. In this treatment, in which the messages have a normative context, the literal meaning of the message seems to have a larger impact on proposers' decisions. Also in the JUST treatment, proposers react averse to very high messages.

In general, the results in Table 2 show, that either messages larger than 5, or best-responses larger than 5, lead proposers to decrease their offer (or at least, not increase the offer further).

4.5 Discussion

Our experimental results show that there is substantial pooling, in our case on the message 5. These messages are considerably inflated. Participants generally have lower rejection thresholds as compared to the literal meaning of the messages. We interpret this as evidence, that low types pool on higher messages in order to avoid smaller offers.[14]

The observed proposer beliefs indicate that communication in our experiment is not cheap-talk. Rather, we observe substantial correlation of proposer beliefs and messages. These beliefs exhibit a common pattern in both treatments: higher messages are believed to come from types with higher minimum acceptance thresholds (tough bargaining types).
In turn, this pattern is in line with responders' beliefs which indicate that sending messages of the lower half of the messages space is associated with lower offers. This can explain why responders do not use the messages of the lower part of the message space. They refrain from sending low messages because if they did, the proposers would believe that the responder is a weak bargaining type which can be exploited by sending a low offer. This observation is in line with our theoretical analysis, that low-type responders do not reveal their true preferences. The shape and correspondence of both responder and proposer beliefs suggests that the lack of low messages in both treatments is meaningful and not the arbitrary result of of communication being cheap talk.

When it comes to high messages, we observe two surprising results, also with respect to theory. First, proposers' beliefs in response to high messages indicate that these messages come from types which cannot exists according to the assumptions of the Fehr-Schmidt model. For instance, suppose a proposer beliefs that a certain message was sent by a type with a minimum acceptance threshold of seven out of ten points. This contradicts the Fehr-Schmidt model according to which any responder should at least accept an offer of five out of ten points, independent of α and β.[15] Nevertheless we observe many cases in which proposers' expect the minimum acceptance threshold to be seven or higher with positive probability.
These beliefs can give an explanation, why proposers frequently send lower offers than their

[14] An alternative explanation for the pooling on 5 is that—apart from fairness considerations—the equal split of the points can serve as a focal point. Given that there are multiple equilibria, our participants might "coordinate" on this focal point although preferences are (at least to some extend) misaligned (Mehta, Starmer, & Sugden, 1994a and 1994b).

[15] Such a rejection threshold would also be impossible according to the original parametrization of the model by Levine (1998)

best response for high messages. Beliefs, that an offer of 5 is rejected, leaves proposers in the behind-averse domain of the Fehr-Schmidt model. Hence a sufficiently behind averse proposer may maximize her utility, by sending an offer which is less than her best response, if the best response is larger than 5. In the light of the Fehr-Schmidt model, offering less than the best response must come from beliefs in types who reject an offer of 5 or more. Otherwise, the proposer would (instead of the behind averse domain) be in the ahead averse domain. In this domain, we would exclusively expect offers larger than the (selfish) best response.
The deviation from proposers risk-neutral selfish best responses can hence be explained by qualitative predictions of the Fehr-Schmidt model. Although proposers believe in high rejection thresholds, they do not frequently send larger offers than the equal split, because of inequality aversion.

The second surprising finding is that responders' do only rarely send high messages. While the above arguments might explain why responders have the (correct) expectation that offers do not increase much beyond a message of five, they nevertheless believe in an increase (according to the belief reports). However, also responders deviate significantly downwards from their best-response message. This is also a puzzle for the Fehr-Schmidt model. In order to explain that responders do no use very high messages in the context of the Fehr-Schmidt model, we would need to assume that the vast majority of responders is ahead averse ($\beta > 1/2$). This assumption contradicts the frequently observed behavior of proposers, who are willing to implement advantageous inequality whenever possible. Hence we would also expect a substantial fraction of non-ahead averse responders to use the messages of the upper half of the message space.
To wrap up, our results indicate that responder communication in the ultimatum game is not cheap-talk. The messages influence proposers' beliefs about the social preference type of the responder. In this sense, there is information transmission through communication. However, we can also show that — in line with our theoretical result — responders do not fully reveal their types. There is hence a limit to information transmission.

4.6 Conclusion

Our experimental results show, that there are indeed limits to information transmission in the ultimatum game with communication. There is substantial pooling on messages around 5, but actual rejections thresholds are smaller than 5 very often. However, sending lower (*i.e. truthful*)

messages would often lead to proposers sending lower offers. Messages larger than 5 are to some degree believed to come from very extreme types. Yet, offers do not substantially increase for these high messages. In line with the model of inequality aversion, most proposers do not send more than 50% of the pie and some even react aversely to very high messages.

Our results show that communication, which is cheap talk according to standard theory, can have large impacts on behavior in bargaining situations. Messages can contain information about social preferences, for example about inequality aversion. Although there are limits to information transmission, the transfered information about preferences influences behavior.

References

Andersson, O., Galizzi, M. M., Hoppe, T., Kranz, S., Van Der Wiel, K., & Wengström, E. (2010). Persuasion in experimental ultimatum games. *Economics Letters*, 108(1), 16-18.

Boles, T. L., Croson, R. T., & Murnighan, J. K. (2000). Deception and retribution in repeated ultimatum bargaining. *Organizational Behavior and Human Decision Processes*, 83(2), 235-259.

Croson, R., Boles, T., & Murnighan, J. K. (2003). Cheap talk in bargaining experiments: lying and threats in ultimatum games. *Journal of Economic Behavior & Organization*, 51(2), 143-159.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817-868.

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

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

Güth, W., & Kocher, M. G. (2014). More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior & Organization*, 108, 396-409.

Hossain, T., & Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies*, 80(3), 984-1001.

Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1(3), 593-622.

Lundquist, T., Ellingsen, T., Gribbe, E., & Johannesson, M. (2009). The aversion to lying. *Journal of Economic Behavior & Organization*, 70(1), 81-92.

McKelvey, R. D., & Page, T. (1990). Public and private information: An experimental study of information pooling. *Econometrica: Journal of the Econometric Society*, 1321-1339.

Mehta, J., Starmer, C., & Sugden, R. (1994a). The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, 84(3), 658-673.

Mehta, J., Starmer, C., & Sugden, R. (1994b). Focal points in pure coordination games: An experimental investigation. *Theory and Decision*, 36(2), 163-185.

Oosterbeek, H., Sloof, R., & Van De Kuilen, G. (2004). Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7(2), 171-188.

Rankin, F. W. (2003). Communication in ultimatum games. *Economics Letters*, 81(2), 267-271.

Roth, A. (1995). Bargaining Experiments. In J. Kagel & A. Roth (Eds.) *Handbook of Experimental Economics*. Princeton University Press.

Selten, R. (1967). Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopolexperiments. In Sauerman, H.(ed.) Beiträge zur Experimentellen Wirtschaftsforschung, Tübingen: JCB Mohr (Paul Siebeck).

Sheremeta, R. M., & Shields, T. W. (2013). Do liars believe? Beliefs and other-regarding preferences in sender–receiver games. *Journal of Economic Behavior & Organization*, 94, 268-27

Zultan, R. I. (2012). Strategic and social pre-play communication in the ultimatum game. *Journal of Economic Psychology*, 33(3), 425-434.

Appendix

4.A Figures and Tables

Choice	Option 1	Option 2	
		Safe choice	Risky choice
# 1	1 MU	3 MU with probability 0/20	0 MU with probability 20/20
# 2	1 MU	3 MU with probability 1/20	0 MU with probability 19/20
# 3	1 MU	3 MU with probability 2/20	0 MU with probability 18/20
# 4	1 MU	3 MU with probability 3/20	0 MU with probability 17/20
# 5	1 MU	3 MU with probability 4/20	0 MU with probability 16/20
# 6	1 MU	3 MU with probability 5/20	0 MU with probability 15/20
# 7	1 MU	3 MU with probability 6/20	0 MU with probability 14/20
# 8	1 MU	3 MU with probability 7/20	0 MU with probability 13/20
# 9	1 MU	3 MU with probability 8/20	0 MU with probability 12/20
# 10	1 MU	3 MU with probability 9/20	0 MU with probability 11/20
# 11	1 MU	3 MU with probability 10/20	0 MU with probability 10/20
# 12	1 MU	3 MU with probability 11/20	0 MU with probability 9/20
# 13	1 MU	3 MU with probability 12/20	0 MU with probability 8/20
# 14	1 MU	3 MU with probability 13/20	0 MU with probability 7/20
# 15	1 MU	3 MU with probability 14/20	0 MU with probability 6/20

Table A1: Elicitation of risk-preferences. Subjects choose between a certain payoff of 1 MU (Option 1) or a risky payoff of either 3 MU or 0 MU (Option 2). Risk neutrality requests choosing Option 1 in choices # 1 to 7. A higher cutoff corresponds to higher risk aversion. We estimate participants cutoffs by a maximum-likelihood estimation, including a tremble parameter accounting for multiple switching points.

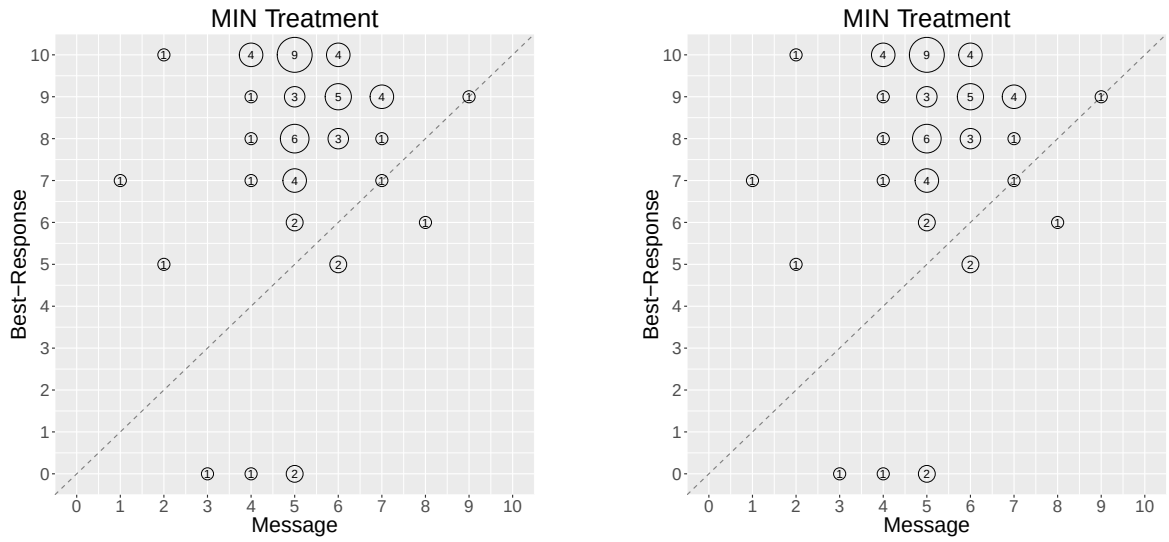

Figure A1: Bubble plots of messages and best-responses. Bubble sizes and numbers indicate the observed frequency of the respective Message - Best-Response combination.

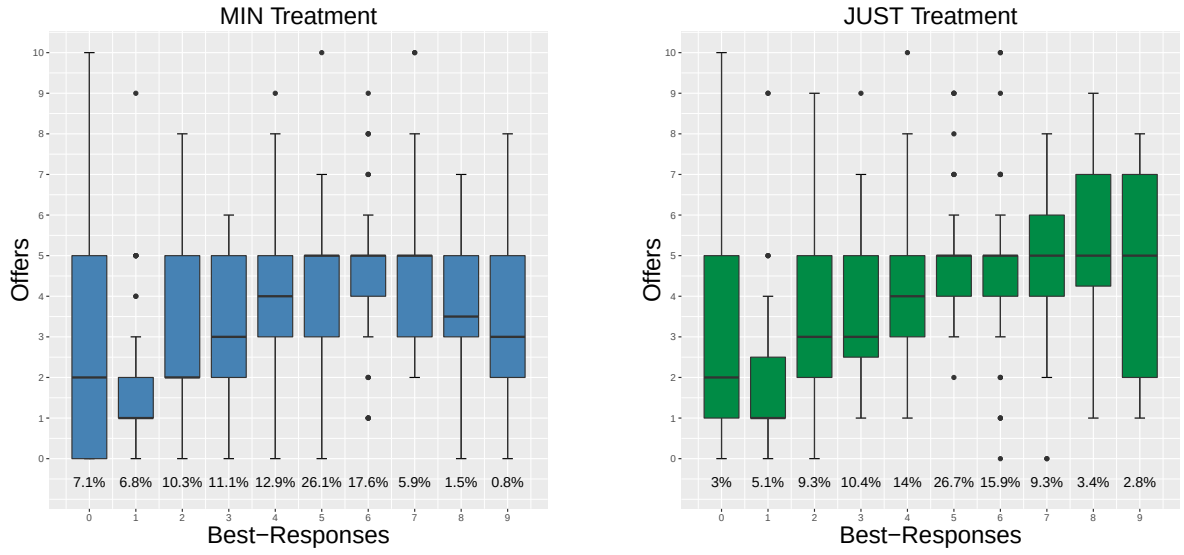

Figure A2: Boxplots of offers conditional on best-responses. Relative frequencies of best-responses are indicated below boxplots.

Offer - Best-Response	Treatment	
	MIN	JUST
Message	-0.110**	-0.161***
	(0.048)	(0.054)
Constant	-1.989	-2.550
	(1.976)	(2.891)
Clusters	60	48
R^2	0.0586	0.1057

Table A2: Linear regressions of the difference between offer and best-response on messages. Standard errors in parentheses are clustered on the participant level. Asterisks: *** p<0.01, ** p<0.05. Additional controls in all models: *male*, *age*, *available money per month*, and a self reported *reliability-of-answers* measure. We further control for the quality of our risk preference estimation by adding the tremble parameter and the standard deviation of the likelihood function.

4.B Communication in the ultimatum game with Fehr-Schmidt utility

In the ultimatum game, two players bargain over the distribution of a fixed amount of monetary units (MU) which we normalize to 1. The proposer offers a share $s \in [0, 1]$ of the MU to the responder and keeps the rest $(1 - s)$ for himself. The responder can accept or reject the offer. If the responder accepts, the players receive the proposed shares as payoffs. If the responder rejects, both players receive a payoff of zero.

Assume a utility function of the Fehr-Schmidt model of inequality aversion (Fehr & Schmidt, 1999). The utility function of both player types is:

$$u_i(\pi_i, \pi_j) = \pi_i - \alpha_i \, max\{0, \pi_j - \pi_i\} - \beta_i \, max\{0, \pi_i - \pi_j\}.$$

Here, π_i is the income of player i , $\alpha_i \in [0, \infty)$ measures the aversion to disadvantageous inequality and $\beta_i \in [0, 1)$ is the aversion to disadvantageous inequality of player i with $\alpha_i \geq \beta_i$. Let (α_R, β_R) denote the responder's preference parameters and (α_P, β_P) the proposer's parameters. Throughout the proof, we will make use of the original Proposition 1 from Fehr & Schmidt (1999, page 826). The following three statements are due to Fehr & Schmidt (1999), e reproduce them here for convenience of the reader.

First, in the Fehr-Schmidt model, the equal split ($s = 0.5$) always yields positive utility to both players and is always accepted by responders. A proposer will never offer more than the equal split because all offers with $s > 0.5$ yield lower utility than the equal split (which is accepted with certainty).

Second, α_R defines a minimum acceptance threshold for the responder. All offers below this

threshold yield negative utility for the responder and are hence rejected.[16] Since no proposer offers more than the equal split ($s \leq 0.5$), the acceptance threshold $\hat{s}_R$ is located in the domain of being behind and is defined by the offer that yields zero utility to the responder for a given α_R:

$$\begin{aligned} u_R(\hat{s}_R, 1-\hat{s}_R) = \hat{s}_R - \alpha_R(1-\hat{s}_R s) &= 0 \\ \hat{s}_R &= \frac{\alpha_R}{1+2\alpha_R} \end{aligned} \tag{4.1}$$

With $\hat{s}_R \in [0, 0.5)$ because $\alpha_R \in [0, \infty)$. All offered shares which are smaller than $\hat{s}_R$ are rejected, and all offer larger than $\hat{s}_R$ are accepted. We assume that offers which yield zero utility to the responder are accepted.

Third, if the proposer *knows* the type of the responder, equilibrium offers fulfill:

$$s^* = \begin{cases} 0.5, & \text{if } \beta_P > \frac{1}{2} \\ s', s' \in [\hat{s}_R, 0.5] & \text{if } \beta_P = \frac{1}{2} \\ \hat{s}_R & \text{if } \beta_P < \frac{1}{2} \end{cases} \tag{4.2}$$

This means that if the proposer knows the type of the responder, ahead averse proposers ($\beta_P > \frac{1}{2}$) will offer the equal split, while proposers with limited ahead aversion ($\beta_P < \frac{1}{2}$) will offer the smallest amount which will be accepted (and proposers with $\beta_P = \frac{1}{2}$ are indifferent between the options).

Communication

We now introduce communication to the ultimatum game. Before the proposer makes an offer, the responder sends a message $m \in [0, 1]$ to the proposer. This message may contain the minimum share that must be offered for the responder to accept the offer. However, the message is non-binding and costless and hence need not contain the true minimum share $\hat{s}_R$. Only after receiving the message, the proposer decides on an offer. The message hence potentially influences the proposer's beliefs about the responders social preference type.

We are interested only in showing that full separation of types is impossible in equilibrium. We start by determining the proposer reaction to a revealed responder.

We then assume that there exist different responder types and that these types fully separate. We next determine the proposers' reactions to full separation and then show that responders have an incentive to deviate from full separation.

[16] Technically, α_P defines a minimum acceptance threshold also for proposers.

Suppose some responder type truthfully reveals her acceptance threshold $\hat{s}_R$ to the proposer via communication. This means that after receiving the message, the proposer's beliefs are such that he knows the type (the α_R) of the responder.

Given these beliefs, every proposer type maximizes his utility according to equation (4.2). For ahead averse proposers with $\beta_P > \frac{1}{2}$ the revealed information about the responder's type does not matter because they offer the equal split ($s^* = 0.5$) anyway. However, proposers with limited ahead aversion ($\beta_P < \frac{1}{2}$) will offer $\hat{s}_R$, the smallest amount which will be accepted by the revealed responder type. Hence, these types react one-to-one to the revealed information.

We now analyze whether responders can reveal their types in equilibrium. Assume that the distribution of behind aversion $F(\alpha)$ is such that there exist at least two different rejection thresholds in the population which are denoted by $\hat{s_i}^{high} > \hat{s_i}^{low}$. We call the responders with the respective thresholds the *high* and *low* types.

Suppose responders send messages such that they fully separate and every responder reveals her type. This means, that proposers' beliefs are such that they know the responders' types. Following from the argument above (and by equation (4.2)), all proposers with $\beta_P > \frac{1}{2}$ will offer the equal split, independent of the responder type. However, the low [high] type will receive an offer of $\hat{s_i}^{low}$ [$\hat{s_i}^{high}$] from proposers with $\beta_P < \frac{1}{2}$.

Different responder-types will receive different offers from the limited ahead averse proposers. Now since $\hat{s_i}^{high} > \hat{s_i}^{low}$, the low type has hence an incentive to deviate from full separation, for example by imitating the message of the high type. A low type responder could increase utility by receiving the higher offer of $\hat{s_i}^{high}$. The low type responders have hence an incentive to send messages that lead proposers to believe that they are of a higher type than they actually are. Because then, proposers will send higher offers.

It is hence impossible that different responder types fully separate with respect to communication in equilibrium. ■

4.C Experimental instructions

The instructions are translated from german. The instructions are for responders. Whenever possible, proposers received identical instructions. Otherwise, the structure and content are the same, but explained from the viewpoint of the proposer.

Instructions

Welcome to the experiment. Please read the following instructions carefully. In case you have any questions now or during the experiment, please raise your hand. Please stop any communication with other participants from now on.

The experiment consists of 3 parts. Part 1 & 2 belong together. Part 3 is *independent* from the first two parts and will also be payed independently. You receive the instructions for part 3 on screen and after parts 1 & 2 are finished.
You can earn points in all three parts of the experiment. At the end of the experiment, *one* part of part 1 & 2 will be randomly selected and payed. The points you earn in part 3 are always payed. The total number of points you earned will be converted to EURO according to the following exchange rate:

1 point = 150 Euro-cents

After all parts of the experiment are completed, there will be a short questionnaire. You receive an additional 2 Euro for completion of the questionnaire. You receive your payment at the end of the experiment in cash under full privacy. No other participant will know, how much money you earned.

Your participant role

There are two different types of participants, type A and B. The types are randomly assigned at the beginning of the experiment and do not change during the experiment. You are type A. In the first two parts, you will be anonymously and randomly matched to one other participant in this room. This participant is of type B.

Part 1

Overview

In the first part of the experiment, you and participant B jointly have 10 points at your disposal. Together with participant B, you will decide how to distribute the 10 points among you. Participant B has to offer you an integer amount between 0 and 10 Points. You can either accept or reject this offer. If you accept the offer, the 10 points will be exactly distributed as proposed by participant B. This means, you receive thze offered amount of points and participant B receives the rest. If you reject the offer, you and participant B will both receive 0 points.

Procedure of part 1

In the beginning, you decide on two values: A **message** and a **minimum**.

With the **message**, you can tell participant B, **what the smallest offer is, you are barely willing to accept.**

Additionally, you will decide on a **minimum**. This minimum, is the **actual minimum number of points, which you are barely willing to accept.** The minimum is later used to determine your payoff. Participant B does not know, which minimum you chose. Independently of which message you chose above, you can enter any minimum. Participant B does also know all of this.

Participant B does not receive your message directly. Instead, he decides on an offer, conditional on every of the eleven possible messages (0, 1, 2,..., 10) he might receive. This means, for every possible message, he makes an offer of how many points (from 0 to 10, integer) you receive in this case and how many (the rest) he receives.

At the end, you will receive the specific offer, participant B made for the message you sent. If this offer is equal or larger than your selected minimum, the offer is automatically accepted and the distribution of points which was submitted by participant B is implemented. If the offer is smaller than your selected minimum, is is automatically rejected and you and participant B do not receive any points.

Part 2

In the second part of the experiment we want to know about your beliefs how participant B is going to decide. More precisely, we want to know your assessment of the following question: Assume, you sent the message Y to participant B. **What is the probability of receiving an offer of X in this case?**

As described above, participant B chooses an offer for every message he might receive. For every of the eleven possible messages, you hence have to make a percentage estimate how likely participant B is to decide on a specific offer. To this end, you will see a diagram for each message. In this diagram, you can enter your probability assessment on how likely participant B is to make an specific offer. In the figure you see such a diagram as an example.

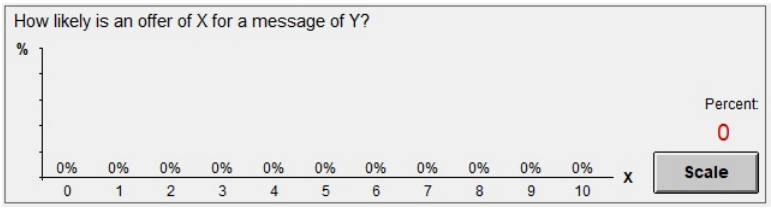

You will hence enter for each possible message, how likely you think it is to receive a specific offer. The precise handling of the diagram is explained later on screen.

In part 2, you can either receive 7 or 3 points. Your chance of receiving 7 points increases in the precision of your probability assessment. The precision of your assessment increases, the more it is in line with the actual decision behavior of participant B. At the end of the experiment, one of the 11 diagrams is randomly selected and payed. How we compute your payoff which rewards your precision is explained in the appendix. If you are not interested in the details, you may ignore the instructions without any doubt.

It is important for you to know, that the chance of receiving a high payoff is maximal, if you assess the behavior of participant B correctly. It is our intention, that you have an incentive to think carefully about the behavior of participant B. We want, that you are rewarded if you have assessed the behavior well and made a respective report.

Payment of the probability assessments

At the end of your assessment, you will receive the 7 points with a certain chance (p) and with $(1-p)$, you receive 3 points. You can influence your chance p with your assessment in the following way:

As described above, you will —for a given message— make percentage estimates how likely each possible offer is to be selected by participant B. One of these offers is the actually selected. Your chance p is initially set to 1 (hence 100%). However, there will be deductions, if your assessments are wrong. The deductions are first squared and then divided by two.

For example, if you place 70% on a specific offer, but participant B selects another offer, your deviation is equal to 0.7. Hence, we deduct $0.70 * 0.70 * \frac{1}{2} = 0.245$ (ca. 25%) from p.
For the offer, which is actually selected by participant B, it is bad if your assessment is far away from 100%. Again, your deviation is squared, halved and deducted. For example if you only place 60% probability on the actually selected offer, we will deduct $(1-0.60) * (1-0.60) * \frac{1}{2} = 0.40 * 0.40 * \frac{1}{2} = 0.08$ (8%) from p.

At the end, all deductions are summed up and the smaller the sum of squared deviations is, the better was your assessment. For those who are interested, we show the mathematical formula according to which we compute the quality of your assessment and hence your chance p of receiving 7 points.

$$p = 1 - \frac{1}{2}\left[\sum (q_{wrong,i})^2 + (1 - q_{correct})^2\right]$$

The value of p of your assessment will be computed and displayed to you at the end of the experiment. The higher p is, the better your assessment was and the higher your chance to receive 7 points (instead of 3) in this part. At the end of the experiment, the computer will draw a random number between 0 and 100. If this number is smaller or equal to p, you receive 7 points. If the number is larger than p you receive 3 points.

Summary

In order to have a high chance to receive the large payment, it is your aim to achieve as few deductions from p as possible. This works best, if you have an good assessment of the behavior of participant B and report that assessment truthfully.

If you are ready or have any questions, please raise your hand.

Abgrenzung

Die erste Studie *"Biases in Beliefs: Experimental Evidence"* habe ich zusammen mit Irenaeus Wolff verfasst. Die Idee und das experimentelle Design wurden gemeinsam erarbeitet. Die Programmierung, die Durchführung im Labor und die Datenanalyse stammen von mir. Das Papier wurde von uns beiden im ständigen Austausch gemeinsam geschrieben.
Das zweite Papier *"Belief Uncertainty and Stochastic Choice"* ist ebenfalls eine Zusammenarbeit mit Irenaeus Wolff. Auch in dieser Studie stammen die Idee und das Design aus unzähligen gemeinsamen Gesprächen und einem Seminar. Die Programmierung, die Durchführung im Labor und die Datenanalyse stammen von mir. Das Papier wurde von uns beiden im ständigen Austausch gemeinsam geschrieben.
Die dritte Studie *"Responder Communication in Ultimatum Bargaining"* habe ich zusammen mit Fabian Dvořák verfasst. Die Idee und das experimentelle Design wurden gemeinsam erarbeitet. Die Progammierung stammt von Fabian und ich habe das Experiment im Labor durchgeführt. Die Datananalyse stammt weitestgehend von mir, wobei jedoch die Schätzung der Risikoparameter von Fabian durchgeführt wurde. Das Papier wurde von uns beiden im ständigen Austausch gemeinsam geschrieben.

Bibliography

Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80(1), 123-129.

al-Nowaihi, A., & Dhami, S. (2015). Evidential Equilibria: Heuristics and Biases in Static Games of Complete Information. *Games*, 6(4), 637-676.

Alós-Ferrer, C., Granić, D. G., Shi, F., & Wagner, A. K. (2012). Choices and preferences: Evidence from implicit choices and response times. *Journal of Experimental Social Psychology*, 48(6), 1336-1342.

Alós-Ferrer, C., Granić, D. G., Kern, J., & Wagner, A. K. (2016). Preference reversals: Time and again. *Journal of Risk and Uncertainty*, 52(1), 65-97.

Armantier, O., & Treich, N. (2013). Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. *European Economic Review*, 62, 17-40.

Andersson, O., Galizzi, M. M., Hoppe, T., Kranz, S., Van Der Wiel, K., & Wengström, E. (2010). Persuasion in experimental ultimatum games. *Economics Letters*, 108(1), 16-18.

Babad, E., & Katz, Y. (1991). Wishful thinking—against all odds. *Journal of Applied Social Psychology*, 21(23), 1921-1938.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593-1636.

Bauer, D. & Wolff, I. (2017a). Biases in Beliefs: Experimental Evidence. *Mimeo*.

Bauer, D., & Wolff, I. (2017b). Belief uncertainty and stochastic choice. *Mimeo*.

Bar-Hillel, M., & Budescu, D. V. (1995). The elusive wishful thinking effect. Thinking & Reasoning, 1(1), 71-103.

Bar-Hillel, M., Budescu, D. V., & Amar, M. (2008). Predicting World Cup results: Do goals seem more likely when they pay off? *Psychonomic Bulletin & Review*, 15(2), 278-283.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1963). Stochastic models of choice behavior. *Behavioral Science*, 8(1), 41-55.

Bellemare, C., Kröger, S., & Van Soest, A. (2008). Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities. *Econometrica*, 76(4), 815-839.

Blanco, M., Engelmann, D., Koch, A. K., & Normann, H. T. (2010). Belief elicitation in experiments: is there a hedging problem?. *Experimental Economics*, 13(4), 412-438.

Blanco, M., Engelmann, D., Koch, A. K., & Normann, H. T. (2014). Preferences and beliefs in a sequential social dilemma: a within-subjects analysis. *Games and Economic Behavior*, 87, 122-135.

Boles, T. L., Croson, R. T., & Murnighan, J. K. (2000). Deception and retribution in repeated ultimatum bargaining. *Organizational Behavior and Human Decision Processes*, 83(2), 235-259.

Breitmoser, Y. (2015). Knowing me, imagining you: Projection and overbidding in auctions. *Working paper*, accessed 2017/09/06, https://mpra.ub.uni-muenchen.de/62052/

Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity Analysis of Production and Allocation*, 13(1), 374-376.

Camerer, C., & Hua Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4), 827-874.

Camerer, C., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *The American Economic Review*, 89(1), 306-318.

Chabris, C. F., Morris, C. L., Taubinsky, D., Laibson, D., & Schuldt, J. P. (2009). The Allocation of Time in Decision-Making. *Journal of the European Economic Association*, 7(2/3), 628-637.

Charness, G., & Grosskopf, B. (2001). Relative payoffs and happiness: an experimental study. *Journal of Economic Behavior & Organization*, 45(3), 301-328.

Charness, G., & Levin, D. (2005). When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect. *The American Economic Review*, 95(4), 1300-1309.

Christensen-Szalanski, J. J., & Willham, C. F. (1991). The hindsight bias: A meta-analysis. *Organizational Behavior and Human Decision Processes*, 48(1), 147-168.

Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies*, 75(3), 729-762.

Critcher, C. R., & Dunning, D. (2013). Predicting persons' versus a person's goodness: Behavioral forecasts diverge for individuals versus populations. *Journal of Personality and Social Psychology*, 104(1), 28.

Critcher, C. R., & Dunning, D. (2014). Thinking about Others versus Another: Three Reasons Judgments about Collectives and Individuals Differ. *Social and Personality Psychology Compass*, 8(12), 687-698.

Croson, R., Boles, T., & Murnighan, J. K. (2003). Cheap talk in bargaining experiments: lying and threats in ultimatum games. *Journal of Economic Behavior & Organization*, 51(2), 143-159.

Cubitt, R. P., Starmer, C. & Sugden, R. (2001). Discovered preferences and the experimental evidence of violations of expected utility theory. *Journal of Economic Methodology*, 8(3), 385-414.

Danz, D. N., Fehr, D., & Kübler, D. (2012). Information and beliefs in a repeated normal-form game. *Experimental Economics*, 15(4), 622-640.

Danz, D. N., Madarász, K., & Wang, S. W. (2014). The Biases of Others: Anticipating Informational Projection in an Agency Setting. *Working Paper*, accessed 2017/06/06, http://works.bepress.com/kristof_madarasz/42/

Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment?. *Organizational Behavior and Human Decision Processes*, 65(3), 201-211.

Delavande, A., Giné, X., & McKenzie, D. (2011a). Eliciting probabilistic expectations with visual aids in developing countries: how sensitive are answers to variations in elicitation design? *Journal of Applied Econometrics*, 26(3), 479-497.

Delavande, A., Giné, X., & McKenzie, D. (2011b). Measuring subjective expectations in developing countries: A critical review and new evidence. *Journal of Development Economics*, 94(2), 151-163.

Dhami, S. (2016). *The Foundations of Behavioral Economic Analysis*. Oxford University Press, Oxford, UK.

Dickhaut, J., Smith, V., Xin, B., & Rustichini, A. (2013). Human economic choice as costly information processing. *Journal of Economic Behavior & Organization*, 94, 206-221.

Dow, J., & da Costa Werlang, S. R. (1992). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60(1), 197-204.

Engelberg, J., Manski, C. F., & Williams, J. (2011). Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition. *Journal of Applied Econometrics*, 26(7), 1059-1078.

Engelmann, D., & Strobel, M. (2012). Deconstruction and reconstruction of an anomaly. *Games and Economic Behavior*, 76(2), 678-689.

Ellingsen, T., Johannesson, M., Tjøtta, S., & Torsvik, G. (2010). Testing guilt aversion. *Games and Economic Behavior*, 68(1), 95-107.

Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology*, 87(3), 327.

Epley, N., & Gilovich, T. (2016). The mechanics of motivated reasoning. *The Journal of Economic Perspectives*, 30(3), 133-140.

Etner, J., Jeleva, M., & Tallon, J. M. (2012). Decision theory under ambiguity. *Journal of Economic Surveys*, 26(2), 234-270.

Eyster, E. (2002). Rationalizing the past: A taste for consistency. *Working paper*, accessed 2017/06/06, http://www.lse.ac.uk/economics/people/facultyPersonalPages/facultyFiles/ErikEyster/RationalisingThePastATasteForConsistency.pdf

Falk, A., & Zimmermann, F. (2013). A taste for consistency and survey response behavior. CESifo Economic Studies, 59(1), 181-193.

Fechner, G. (1860/ 1966). Elements of Psychophysics, Vol. 1, New York: Holt, Rinehart and Winston

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817-868.

Festinger, L. (1957). A theory of cognitive dissonance. Stanford, CA: Stanford University Press

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

Fischhoff, B. (1975). Hindsight $\neq$ foresight: the effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288-299.

Fudenberg, D., & Levine, D. K. (2016). Whither game theory? Towards a theory of learning in games. *The Journal of Economic Perspectives*, 30(4), 151-169.

Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MIT press.

Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 141-153.

Gilboa, I., Postlewaite, A. W., & Schmeidler, D. (2008). Probability and uncertainty in economic modeling. *The Journal of Economic Perspectives*, 22(3), 173-188.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. New York, Cambridge university press.

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

Guilbault, R. L., Bryant, F. B., Brockway, J. H., & Posavac, E. J. (2004). A meta-analysis of research on hindsight bias. *Basic and Applied Social Psychology*, 26(2-3), 103-117.

Güth, W., & Kocher, M. G. (2014). More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior & Organization*, 108, 396-409.

Grimm, V., & Mengel, F. (2012). An experiment on learning in a multiple games environment. *Journal of Economic Theory*, 147(6), 2220-2259.

Groeneveld, R. A., & Meeden, G. (1977). The mode, median, and mean inequality. *The American Statistician*, 31(3), 120-121.

Guiso, L., & Parigi, G. (1999). Investment and demand uncertainty. *The Quarterly Journal of Economics*, 114(1), 185-227.

Harless, D. W., & Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica*, 62(6), 1251-1289.

Harrison, G. W., Martínez-Correa, J., & Swarthout, J. T. (2014). Eliciting subjective probabilities with binary lotteries. *Journal of Economic Behavior & Organization*, 101, 128-140.

Harris, A. J., & Hahn, U. (2011). Unrealistic optimism about future life events: a cautionary note. *Psychological Review*, 118(1), 135.

Helweg-Larsen, M., & Shepperd, J. A. (2001). Do moderators of the optimistic bias affect personal or target risk estimates? A review of the literature. *Personality and Social Psychology Review*, 5(1), 74-95.

Hey, J. D., & Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 62(6), 1291-1326.

Hossain, T., & Okui, R. (2013). The binarized scoring rule. *The Review of Economic Studies*, 80(3), 984-1001.

Hollard, G., Massoni, S., & Vergnaud, J. C. (2016). In search of good probability assessors: an experimental comparison of elicitation rules for confidence judgments. *Theory and Decision*, 80(3), 363-387.

Holt, C. A., & Smith, A. M. (2016). Belief Elicitation with a Synchronized Lottery Choice Menu That Is Invariant to Risk Attitudes. *American Economic Journal: Microeconomics*, 8(1), 110-139

Hyndman, K. B., Terracol, A., & Vaksmann, J. (2013). Beliefs and (in) stability in normal-form games. *Working paper*, accessed 2017/06/14, http://lemma.u-paris2.fr/sites/default/files/concoursMCF/Vaksman.pdf.

Hyndman, K., Ozbay, E. Y., Schotter, A., & Ehrblatt, W. Z. (2012). Convergence: an experimental study of teaching and learning in repeated games. *Journal of the European Economic Association*, 10(3), 573-604.

Iriberri, N., & Rey-Biel, P. (2013). Elicited beliefs and social information in modified dictator games: What do dictators believe other dictators do? *Quantitative Economics*, 4(3), 515-547.

Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2), 603-606.

Khwaja, A., Sloan, F., & Salm, M. (2006). Evidence on preferences and subjective beliefs of risk takers: The case of smokers. *International Journal of Industrial Organization*, 24(4), 667-682.

Konovalov, A. & Krajbich, I. (2017). "Revealed Indifference: Using Response Times to Infer Preferences." Working Paper.

Krizan, Z., & Windschitl, P. D. (2007). The influence of outcome desirability on optimism. *Psychological Bulletin*, 133(1), 95.

Krueger, J. I. (2007). From social projection to social behaviour. *European Review of Social Psychology*, 18, 1-35.

Krueger, J. I. (2013). Social projection as a source of cooperation. *Current Directions in Psychological Science*, 22(4), 289-294.

Larwood, L., & Whittaker, W. (1977). Managerial myopia: Self-serving biases in organizational planning. *Journal of Applied Psychology*, 62(2), 194

Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1(3), 593-622.

Loomes, G., & Sugden, R. (1995). Incorporating a stochastic element into decision theories. *European Economic Review*, 39(3), 641-648.

Lundquist, T., Ellingsen, T., Gribbe, E., & Johannesson, M. (2009). The aversion to lying. *Journal of Economic Behavior & Organization*, 70(1), 81-92.

Madarász, K. (2012). Information projection: Model and applications. *The Review of Economic Studies*, 79(3), 961–985.

Manski, C. F. (2002). Identification of decision rules in experiments on simple games of proposal and response. *European Economic Review*, 46(4), 880-891.

Manski, C. F., & Neri, C. (2013) First- and second-order subjective expectations in strategic decision-making: Experimental evidence. *Games and Economic Behavior*, 81, 232-254.

Marks, G., & Miller, N. (1987). Ten years of research on the false consensus effect: An empirical and theoretical review. *Psychological Bulletin*, 102(1), 72.

McKelvey, R. D., & Page, T. (1990). Public and private information: An experimental study of information pooling. *Econometrica*, 58, 1321-1339.

McKelvey, R. D., & Palfrey, T. R. (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10(1), 6-38.

Mehta, J., Starmer, C., & Sugden, R. (1994a). The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, 84(3), 658-673.

Mehta, J., Starmer, C., & Sugden, R. (1994b). Focal points in pure coordination games: An experimental investigation. *Theory and Decision*, 36(2), 163-185.

Mullen, B., Atkins, J. L., Champion, D. S., Edwards, C., Hardy, D., Story, J. E., & Vanderklok, M. (1985). The false consensus effect: A meta-analysis of 115 hypothesis tests. *Journal of Experimental Social Psychology*, 21(3), 262-283.

Moffatt, P. G. (2005). Stochastic choice and the allocation of cognitive effort. *Experimental Economics*, 8(4), 369-388.

Molnár, A., & Heintz, C. (2016). Beliefs About People's Prosociality Eliciting predictions in dictator games. *Working Paper*, accessed 2017/09/06, http://publications.ceu.edu/sites/default/files/publications/molnar-heintz-beliefs-about-prosociality.pdf

Mosteller, F., & Nogee, P. (1951). An experimental measurement of utility. *Journal of Political Economy*, 59(5), 371-404.

Norton, E. C., Wang, H., & Ai, C. (2004). Computing interaction effects and standard errors in logit and probit models. *Stata Journal*, 4, 154-167.

Nyarko, Y., & Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3), 971-1005.

Oosterbeek, H., Sloof, R., & Van De Kuilen, G. (2004). Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7(2), 171-188.

Palfrey, T. R., & Wang, S. W. (2009). On eliciting beliefs in strategic games. *Journal of Economic Behavior & Organization*, 71(2), 98-109.

Plott, C.R. (1996) 'Rational individual behavior in markets and social choice processes: the discovered preference hypothesis', In K.J. Arrow, E. Colombatto, M. Perlman and C. Schmidt (eds) *The Rational Foundations of Economic Behaviour*, Basingstoke: Macmillan, pp. 225–50.

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscienc*e, 19(3), 366-374.

Proto, E., & Sgroi, D. (2017). Biased beliefs and imperfect information. *Journal of Economic Behavior & Organization*, 136, 186-202.

Rankin, F. W. (2003). Communication in ultimatum games. *Economics Letters*, 81(2), 267-271.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59.

Rey-Biel, P. (2009) Equilibrium play and best response to (stated) beliefs in normal form games, *Games and Economic Behavior*, 65(2), 572-585.

Ross, L., Greene, D., & House, P. (1977). The "false consensus effect": An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13(3), 279-301.

Roth, A. (1995). Bargaining Experiments. In J. Kagel & A. Roth (Eds.) *Handbook of Experimental Economics*. Princeton University Press.

Rubinstein, A., & Salant, Y. (2015). "Isn't everyone like me?": On the presence of self-similarity in strategic interactions. Working paper version of Rubinstein & Salant (2016), accessed 2017/09/06, https://pdfs.semanticscholar.org/34ee/9a1799fcb4c43207136437e3a1e3c3ef25a6.pdf

Rubinstein, A., & Salant, Y. (2016). "Isn't everyone like me?": On the presence of self-similarity in strategic interactions. *Judgment and Decision Making*, 11(2), 168.

Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. *Economica*, 5(17), 61-71.

Savage, L. J. (1954) The Foundations of Statistics. New York: John Wiley and Sons. (Second ed., Dover, 1972).

Selten, R. (1967). Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopolexperiments. In Sauerman, H.(ed.) Beiträge zur Experimentellen Wirtschaftsforschung, Tübingen: JCB Mohr (Paul Siebeck).

Selten, R., & Ockenfels, A. (1998). An experimental solidarity game. *Journal of Economic Behavior & Organization*, 34(4), 517-539.

Selten, R., & Stoecker, R. (1986). End behavior in sequences of finite Prisoner's Dilemma supergames A learning theory approach. *Journal of Economic Behavior & Organization*, 7(1), 47-70.

Schlag, K. H., Tremewan, J., & Van der Weele, J. J. (2015). A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics*, 18(3), 457-490.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3), 571-587.

Schotter, A., & Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, 6(1), 103-128.

Shah, P., Harris, A. J., Bird, G., Catmur, C., & Hahn, U. (2016). A pessimistic view of optimistic belief updating. *Cognitive Psychology*, 90, 71-127.

Sheremeta, R. M., & Shields, T. W. (2013). Do liars believe? Beliefs and other-regarding preferences in sender–receiver games. *Journal of Economic Behavior & Organization*, 94, 268-27

Sutter, M., Czermak, S., & Feri, F. (2013). Strategic sophistication of individuals and teams. Experimental evidence. *European Economic Review*, 64, 395-410.

Svenson, O. (1981). Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica*, 47(2), 143-148.

Trautmann, S. T., & van de Kuilen, G. (2015). Belief elicitation: A horse race among truth serums. *The Economic Journal*, 125(589), 2116-2135.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207-232.

Tversky, A., & Kahneman, D. (1974). Heuristics and biases: Judgment under uncertainty. *Science*, 185, 1124-1130.

Van Der Heijden, E., Nelissen, J., & Potters, J. (2007). Opinions on the tax deductibility of mortgages and the consensus effect. *De Economist*, 155(2), 141-159.

Weber, Roberto A. (2003). 'Learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, 44(1): 134-144.

Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5), 806.

Weinstein, N. D. (1989). Effects of personal experience on self-protective behavior. *Psychological Bulletin*, 105(1), 31.

Wolff, I. (2015). When best-replies are not in equilibrium: understanding cooperative behaviour. *Working paper*,accessed 2017/09/06, http://kops.uni-konstanz.de/handle/123456789/33027

Wolff, I. (2017). Lucky Numbers in Simple Games. *Mimeo.*

Yariv, L. (2005). I'll See It When I Believe It–A Simple Model of Cognitive Consistency. *Working Paper*, accessed 2017/06/06, http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.207.2893

Zultan, R. I. (2012). Strategic and social pre-play communication in the ultimatum game. *Journal of Economic Psychology*, 33(3), 425-434.

Picture Credit

Chapter 2, Figure A1 and Chapter 3, Figure C1. All pictures are under a public domain license (CC0 1.0 Universal)

Item set 18:
[1] http://www.publicdomainpictures.net/view-image.php?image=59016&picture=flower-outline-coloring-page&large=1
[2] https://commons.wikimedia.org/wiki/File:Amaryllis_2_%28PSF%29.png
[3] https://openclipart.org/detail/214395/flower-lineart
[4] http://www.publicdomainpictures.net/view-image.php?image=35307&picture=orchidee

Item set 20:
[1] http://www.publicdomainpictures.net/view-image.php?image=20740&picture=baum
[2] http://www.publicdomainpictures.net/view-image.php?image=22857&picture=der-baum-auf-der-bank
[3] http://www.publicdomainpictures.net/view-image.php?image=9111&picture=baum
[4] http://www.publicdomainpictures.net/view-image.php?image=22068&picture=baum-im-feld