No. 111

THURGAU INSTITUTE
OF ECONOMICS
at the University of Konstanz

Irenaeus Wolff
Dominik Bauer

Elusive Beliefs: Why Uncertainty Leads to Stochastic Choice and Errors

Research Paper Series
Thurgau Institute of Economics
and Department of Economics
at the University of Konstanz

Elusive Beliefs: Why Uncertainty Leads to Stochastic Choice and Errors[§]

Irenaeus Wolff Dominik Bauer

Thurgau Institute of Economics & Graduate School of Decision Sciences

University of Konstanz, Germany

wolff@twi-kreuzlingen.ch dominik.3.bauer@uni.kn

This version: 28th February, 2018

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 what belief to report or base their decision on, and hence, act on momentaneous 'belief realizations'. In an experiment, we exogenously manipulate participants' belief uncertainty. We find support for both predictions. Low belief uncertainty leads to fewer errors and thus, higher earnings, even when controlling for the accuracy of participants' beliefs. Second, under low belief uncertainty, observed best response rates are high and increasing in the amount of information we provide. Conversely, high belief uncertainty leads to lower consistency.

JEL classification: C91, D81, D83

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

1 Introduction

This paper is about the consequences of uncertainty in people's subjective beliefs for their decision-making. To examine these consequences, we propose a new way of thinking about subjective beliefs, and show how this leads to stochastic choice. In particular, we relate people's belief uncertainty to the likelihood of suboptimal choices, and to the degree of consistency we should expect

[§]We thank Fabian Dvořák, Thomas Hattenbach, Georg Weizsäcker, Ian Krajbich, Wieland Müller, Tomasz Strzalecki, Wolfgang Luhan, Alexander K. Wagner, Simon Gächter, Roberto Weber, Nick Netzer, 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. Note: the order of authors in the header does not convey information on individual contributions. The authors have multiple projects together and agreed to alternate between being first and second authors. Contact: Chair of Applied Research in Economics, University of Konstanz, Universitätsstraße 10, D-78464 Konstanz, Germany.

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. At a first glance, it sounds hardly surprising that more uncertainty leads to more errors. But there are three aspects we want to stress. First, belief uncertainty does not imply indifference or near-indifference. An agent may be uncertain as to the probability of an event while clearly favouring a certain action. Second, we are talking about decisions that are suboptimal *ex ante*, not decisions that turn out to be mistakes only *ex post*. And third, there is nothing in standard economic theory that would relate uncertainty to *ex-ante* suboptimal decisions. The only widely recognized connection between uncertainty and the quality of decisions is through the expected costs of errors, as in models involving Fechner-type errors. However, in this paper, we relate uncertainty to *ex-ante* suboptimal decisions *controlling* for the expected cost of making an error.

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.

In this paper, we show empirically that such uncertainty matters. One important realization is that more information on an uncertain event does not always lead to more consistent behaviour. In fact, when new information contradicts a person's prior belief, the uncertainty in that person's posterior belief will increase in the amount of information over a certain intervall. Showing these relationships empirically is important for modelling the stochasticity of choice. As a second contribution, we offer a simple model that accounts for the observed effects. In particular, we offer 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 beliefs are 'elusive': 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

[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).

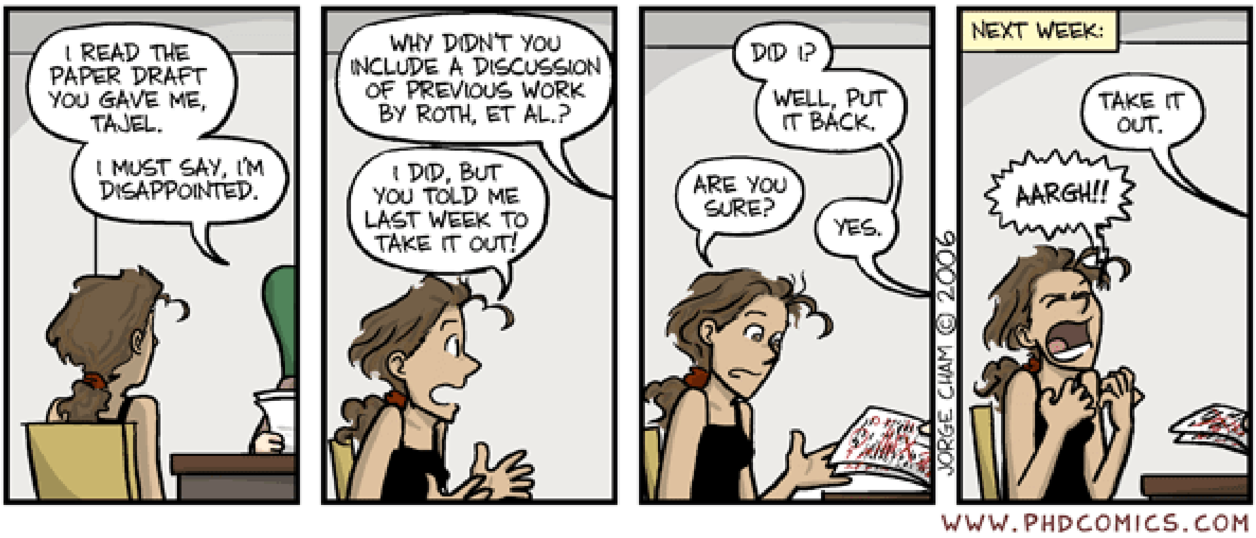

Figure 1: Literary evidence of capricious decision-making by academic supervisors (source: http://phdcomics.com/comics/archive.php?comicid=806, last accessed on 4$^{\text{th}}$ December, 2017.).

time of high uncertainty we should see a higher number of bankruptcies particularly of small firms compared to after a time of little uncertainty.

Our study further will help to understand capricious decision-making by managers and others in similar positions, that may lead to tremendous inefficiencies in the aggregate. The cartoon in Figure 1 provides indirect evidence for the pervasiveness of this phenomenon also in our profession (if the phenomenon was inexistent, the cartoon would not be amusing to anybody). Our model may also contribute to our understanding of what happened before the latest financial crisis: it has been argued that the financial products being traded before the crisis got so complicated that it was impossible for investors to assess the risk the products came with.[3] Consequently, the products were associated with a high degree of uncertainty. However, by our model, high degrees of uncertainty are related not only to a high risk of failure, but also to an increased likelihood of *ex-ante* suboptimal choices; in this perspective, traders not only may have taken excessive risks, but they may have taken too many risks with negative expected value.

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,

[3] See, *e.g.*, Gorton (2009).

from models of belief-based learning like fictitious play or Cournot learning.

Finally, the stochasticity due to belief uncertainty also concerns ourselves as researchers interpreting results from economic experiments. It provides an explanation for the widely varying belief-action-consistency rates reported in the literature.[4] Moreover, it provides an explanation for the observations of Bao, Duffy & Hommes (2013). They test the components of a macroeconomic rational-expectations equilibrium in a cobweb economy. Bao *et al.* find that while forecasts quickly converge to the equilibrium prediction, chosen quantities converge much more slowly. Both the fact that early choices are often not best-responses to forecasts and the fact that choices converge to being best-responses over time are consistent with our story.

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 classic 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

[4] *E.g.*, Costa-Gomes & Weizsäcker (2008) find best-response rates as low as 51% in their 3x3 games, while Rey-Biel (2009) reports 66-69% for similar games. Using the mean squared deviations from the average belief within each game and player role as a proxy for participants' uncertainty (an average of 30% in Costa-Gomes & Weizsäcker *vs.* an average of about 10% in Rey-Biel's study), such a difference is exactly what we would expect.

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.[5] The main question in this literature is how people will make their 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.[6] 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.[7] 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

[5] Even our introductory example in Section 1 is similar to the examples given in this literature, *cf., e.g.*, Gilboa, Postlewaite, and Schmeidler (2008).

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

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

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 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 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.[8]

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.[9] It also would be

[8] 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.

[9] 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.

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.

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:

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

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:

$$E_q[\phi] = \int_0^1 \phi \cdot q(\phi)d\phi = \mu_q \tag{2}$$

In standard theory, the player will then choose $BR(\mu_q)$. In the two-option case we outline here, the critical belief in (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

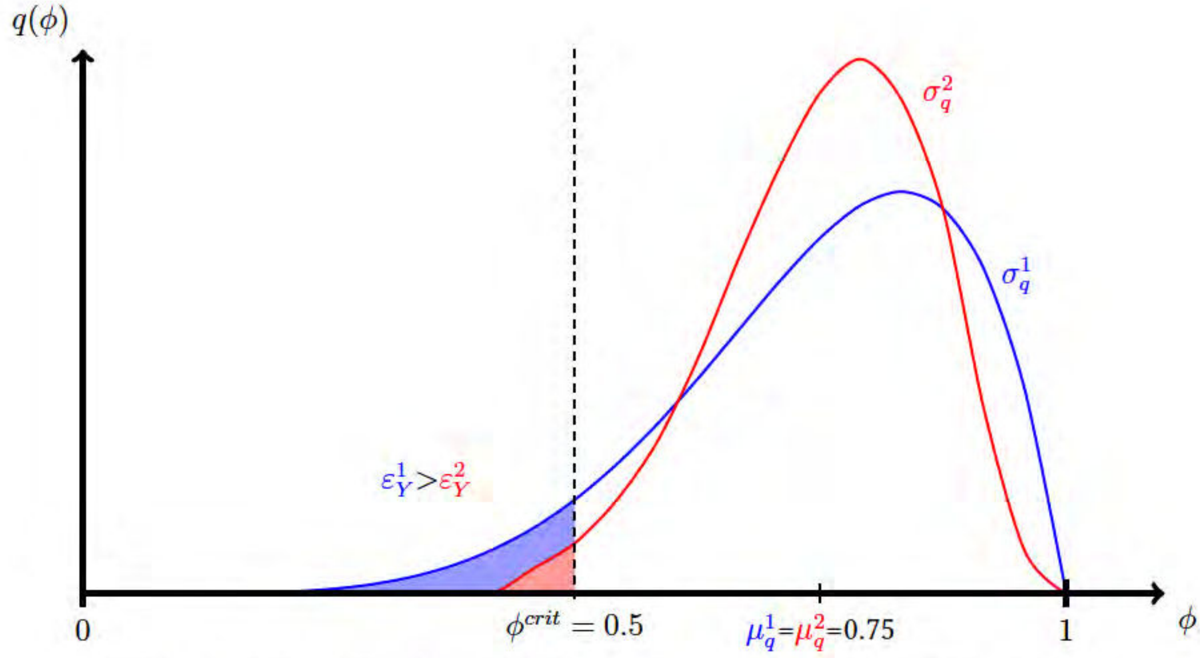

Figure 2: Two belief distributions with identical means but differing variances. The shaded areas indicate the respective error rate ε_Y

equilibrium beliefs because in equilibrium, any action always is a best-response. Throughout the whole paper, we will hence assume $\mu_q \neq 0.5$.

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 2 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 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, 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 an expected-utility difference (ΔEU, where $\Delta EU = |\mu_q - 0.5|$) 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, it is sufficient that *ceteris paribus* the variance of $q(\phi)$ increases. Consider again Figure 2. 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.

Relating our model to 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.

As expected-utility differences (ΔEU) approach zero, the error rate increases *ceteris paribus* in our model. This is also predicted by models in which errors depend directly on 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 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.

Our model thus 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.[10] 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.

[10] If a single draw were to determine both action and belief, we would predict a 100% best-response rate which definitely is rejected by the evidence in the literature as well as in our experiment.

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

$$\widehat{BR} = Prob\Big[BR(\phi^r_A) = BR(\phi^r_B)\Big] = \varepsilon_k^2 + (1-\varepsilon_k)^2 \tag{3}$$

A best response is observed if an error occurs in either none or both of the draws ϕ^r_A, ϕ^r_B. 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.[11]

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 $1 < \alpha < \beta$). ■

[11] 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 2.3.

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 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.

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:

$$p(\phi|n, \alpha, \beta) = \frac{s(\phi|n) \cdot q(\phi; \alpha, \beta)}{t(n, \alpha, \beta)} \tag{4}$$

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:

$$\mu_p = \frac{\alpha + n_X}{(\alpha + n_X) + (\beta + n_Y)} = \underbrace{\frac{\alpha + \beta}{\alpha + \beta + n}}_{1-w} \cdot \underbrace{\frac{\alpha}{\alpha + \beta}}_{\mu_q} + \underbrace{\frac{n}{\alpha + \beta + n}}_{w} \cdot \underbrace{\frac{n_X}{n}}_{\mu_s} \tag{5}$$

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 (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 4, we will use the outlined cases to formulate specific hypotheses for our experiment, which we describe next.

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?"*[12] 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.[13] 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.

[12]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.

[13]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 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 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 (see also Section 4).

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€.

4 Predictions

In this section we specify the hypotheses for our experiment. We base the hypotheses on our model predictions in section 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_B^r, 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_B^r, 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.

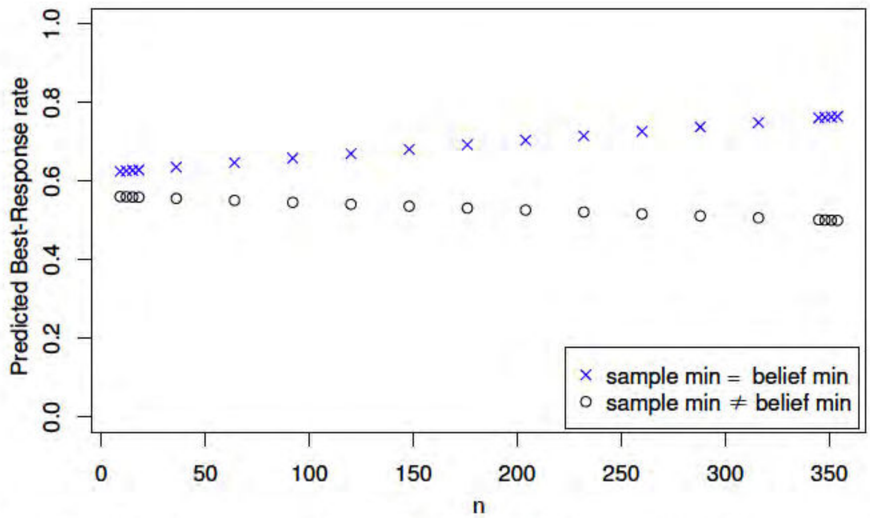

Figure 3: Predicted best-response rates in the four-option game according to our simulation.

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 3. We describe the setup of the simulation in detail in Appendix B.

5 Results

Our most important results are depicted in the left panel of Figure 4, 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 4.

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 decision-specific

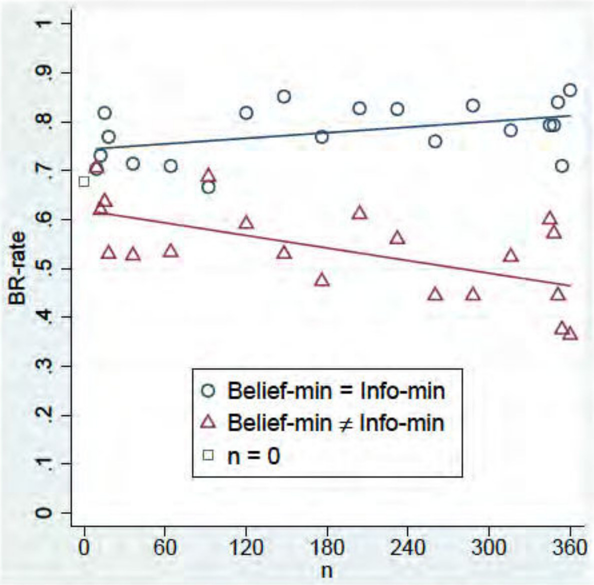

Linear Regression

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 $\neq$ *info-min*	-0.661	0.002

Figure 4: 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).

incentives as a robustness check.

Accounting for error costs and learning

The results in Figure 4 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 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.[14] 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 4). 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^r_B. 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 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

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

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.[15] 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, we 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.[16] 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 making an error. Most importantly, though, the effect of belief uncertainty is robust to adding response times as an alternative measure

[15] 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, $p = 0.312$)

[16] 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.

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

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 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.[17] 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 belief report was

[17] 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	
120	0.10	12.35	$p = 0.020$
354	5.35	16.90	$p = 0.016$

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

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.

6 Conclusion and Discussion

In many cases, people play according to equilibrium only after sufficient experience.[18] 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 interpreting experiment participants' belief-action consistency. Consistency will be predictably low when

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

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 provides an explanation for *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.

Bao, T., Duffy, J., & Hommes, C. (2013). Learning, forecasting and optimizing: An experimental study. *European Economic Review*, 61, 186-204.

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.

Gorton, G. (2009). The subprime panic. *European Financial Management*, 15(1), 10-46.

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.

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

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

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.[19] 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 $E_v[\phi] = \{E[\phi_1], E[\phi_2], E[\phi_3]\}$.

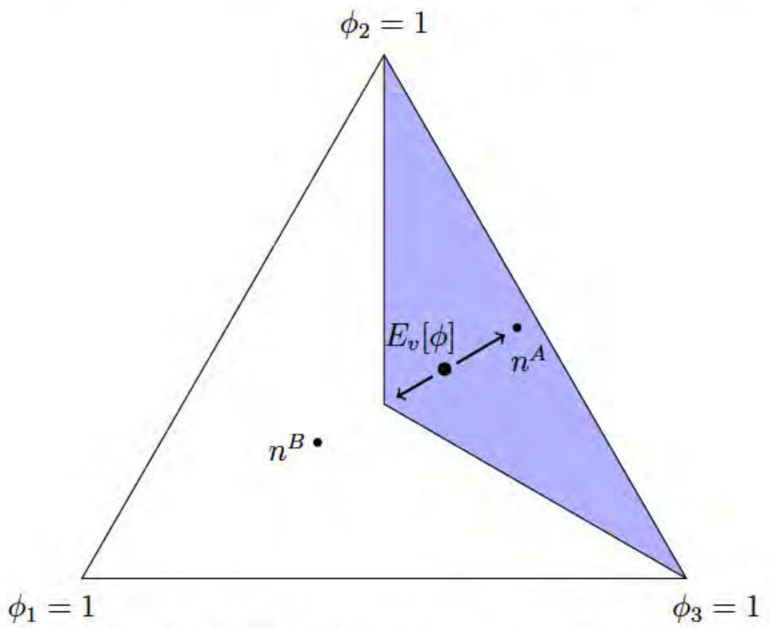

Figure A1: 2-Simplex for all possible beliefs in a three-option game. Exemplary expected choice probabilities and two observed samples n^A and n^B

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 probabilities. Then she commits an error, whenever she draws a ϕ^r that has a different minimum

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

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 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.

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. 19). 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 3 in Section 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.

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

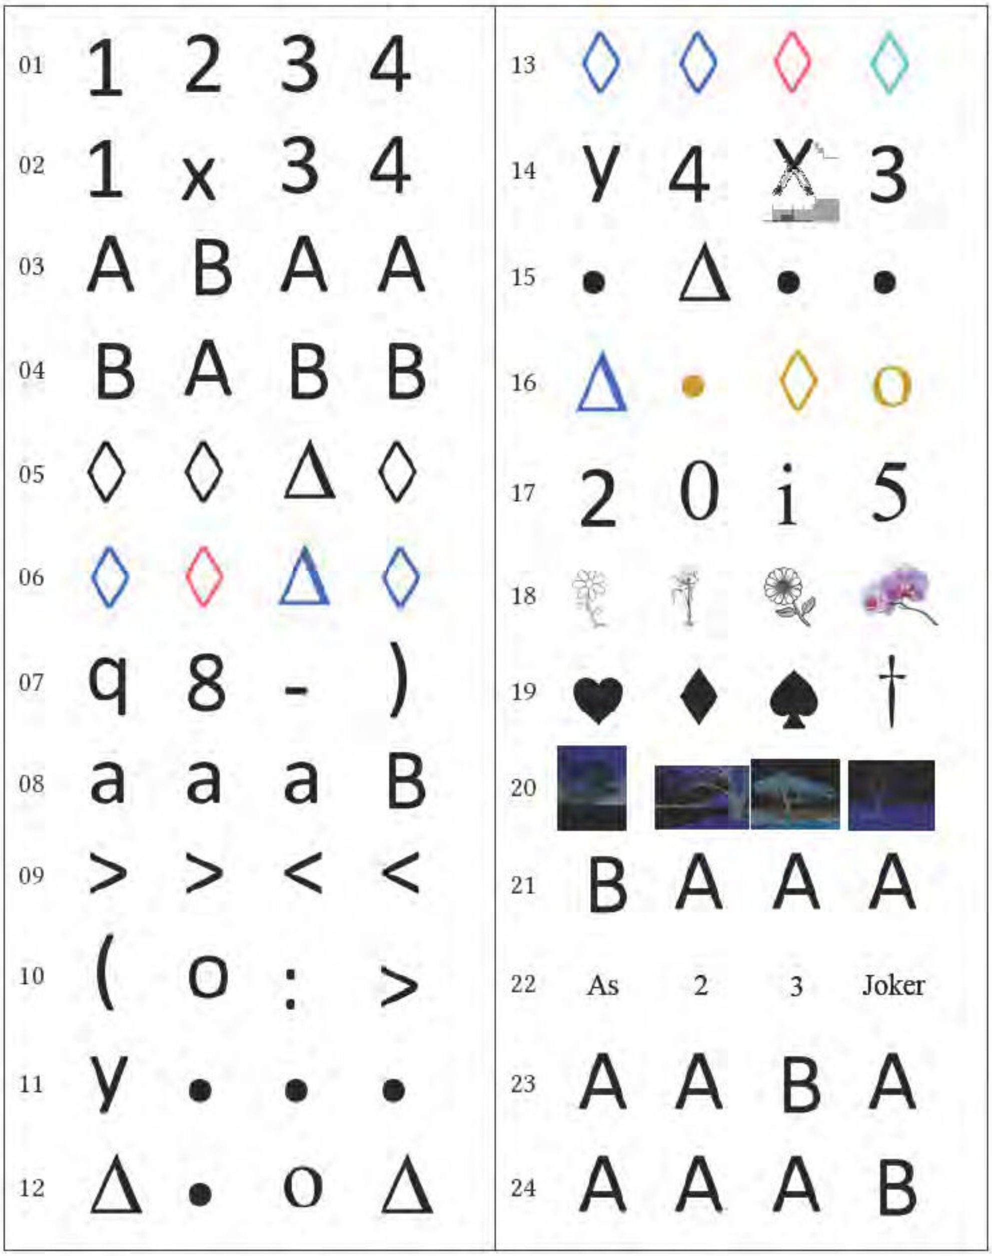

Figure C1: The 24 label sets, used to label the four options of the games. One set for each game.

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.

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.

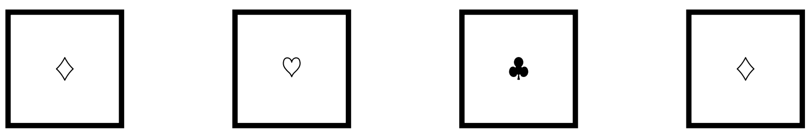

In this example, there are two boxes which are marked with the same symbol. However, the boxes on the most left and most right count as are different boxes.

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.

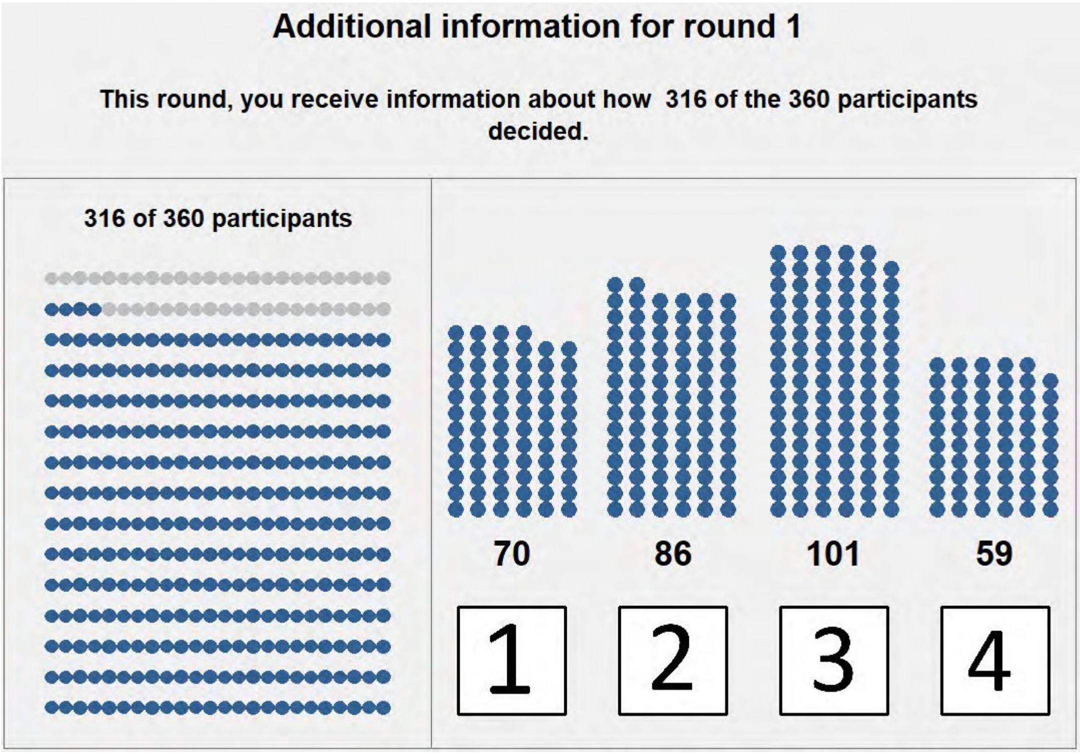

Instructions for decision 2
In decision 2, your payoff also depends on your own decision and on the decision of your matching partner from the earlier experiment. We now explain decision 2 in detail.

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.

THURGAU INSTITUTE
OF ECONOMICS
at the University of Konstanz

Hauptstrasse 90
Postfach
CH-8280 Kreuzlingen 2

T +41 (0)71 677 05 10
F +41 (0)71 677 05 11

info@twi-kreuzlingen.ch
www.twi-kreuzlingen.ch