Behavioral Variation in Tullock Contests*

Aidas Masiliunas[†] Friederike Mengel[‡ §] J. Philipp Reiss[¶]

February 14, 2014

Abstract

We conduct an experiment to uncover the reasons behind the typically large behavioral variation and low explanatory power of Nash equilibrium observed in Tullock contests. In our standard contest treatment, only 7% of choices are consistent with Nash equilibrium which is in line with the literature and roughly what random (uniform) choice would predict (6.25%). We consider a large class of social, risk and some other "non-standard" preferences and show that heterogeneity in preferences *cannot* explain these results. We then systematically vary the complexity of both components of Nash behaviour: (i) the difficulty to form correct beliefs and (ii) the difficulty to formulate best responses. In treatments where *both* the difficulty of forming correct beliefs and of formulating best responses is reduced behavioural variation decreases substantially and the explanatory behaviour of Nash equilibrium increases dramatically (explaining 65% of choices with a further 20% being "close" to NE). Our results show that bounded rationality rather than heterogeneity in preferences is the reason behind the huge behavioral variation typically observed in Tullock contests.

Keywords: rent-seeking, contests, behavioural variation, Nash equilibrium, complexity.

JEL classification: C72, C91, D71, D81

*We would like to thank participants at seminars and conferences in Exeter (Contests, Mechanisms & Experiments), Louvain-la-Neuve, Maastricht (5th MBEES, BEE-Lab meeting), Tilburg (TIBER XI) for comments and suggestions. We would also like to thank David Cooper for sharing their data. A former version of this paper circulated with title '(Strategic) Uncertainty and the Explanatory Power of Nash Equilibrium: An Example using Tullock Contests'.

[†]Department of Economics, Maastricht University, PO BOX 616, 6200MD Maastricht, The Netherlands. e-mail: a.masiliunas@maastrichtuniversity.nl

[‡]University of Essex, Department of Economics, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom. e-mail: fr.mengel@gmail.com

[§]Department of Economics, Maastricht University, PO BOX 616, 6200MD Maastricht, The Netherlands.

[¶]Karlsruhe Institute of Technology (KIT), Institute of Economics (ECON), Schlossbezirk 14, 76131 Karlsruhe, Germany, philipp.reiss@kit.edu

1 Introduction

Many economic, political and social environments can be described as contests in which competing agents expend considerable resources (time, effort, money) in order to increase their chances of winning a “prize”. Examples range from the competition for mates (Andersson and Iwasa (1996)), patents or research grants (Baye and Hoppe (2003)), to promotions or other relative reward schemes in firms (Chen (2003)), lobbying politicians (Baye et al. (1993)), elections (Buchanan and Tullock (1962)), sports competitions (Szymanski (2003)), and ethnic conflicts (Esteban and Ray (2011)). Because of their many applications, these environments have attracted considerable attention in a wide range of fields, both in- and out-side of economics and there is a mature theoretical literature (for a survey see Konrad (2009)).

Experimental economists have tried to understand behaviour in contests empirically. The advantage of conducting experiments on contests (as opposed to empirical field studies) is that effort choices are observable and causal inferences can be drawn via treatment variations. One, maybe surprising, result that has emerged from this literature is that there is huge behavioural variation and Nash equilibrium has little explanatory power in typical Tullock contests (see Millner and Pratt, 1989, Potters et al., 1998, Sheremeta, 2010, among others). Figure 1 illustrates the behavioural variation in experimental Tullock contests and compares it to first- and second-price auctions.[1] The figure shows the cumulative distribution of observed choices relative to the (risk-neutral) Nash prediction for a typical Tullock contest experiment, a first-price auction (FPA), a second-price auction (SPA) and according to NE. According to theory (NE), all the mass should be at 1, because all choices should equal the Nash prediction. Evidently, in the two auction formats (FPA and SPA), the cumulative distribution is pretty similar to theory. Most of the mass is concentrated at 1, where choices equal the Nash prediction. The Nash prediction clearly has something to say about the data here. In the contest, however, this pattern is completely different. Choices seem to have little to do with the unique pure strategy Nash equilibrium and in fact investments are spread across the whole strategy space with no meaningful concentration around any specific value or range of values.

In this paper we try to understand why behavioural variation is so large and why Nash equilibrium (understood more broadly than the risk neutral point prediction) has so little to say about the data in the standard contest. By contrast, most of the literature has focused on the so-called “overbidding phenomenon”, the fact that *on average* investments are above the risk-neutral Nash equilibrium. Explanations have focused on specific preferences or correlates of individual behaviour with specific forms of bounded rationality (see the survey by Sheremeta (2013)). Examples of preference based explanations include spiteful preferences, inequality aversion (Bartling et al. (2009)) or the “joy of winning hypothesis” (Schmitt et al. (2004), Cason et al. (2010)). The “overbidding phenomenon” has also been explained by QRE (Sheremeta (2011), Lim et al. (2014)), distortion of probabilities (Baharad and Nitzan (2008)) or learning

[1] Auctions and contests are both “competitive” allocation games, where everyone bids for a prize but there is one winner. Two differences are that in auctions there is incomplete information, while there is complete information in the contest we study. In addition, in auctions the prize allocation is deterministic, while in the typical contest it is non-deterministic.

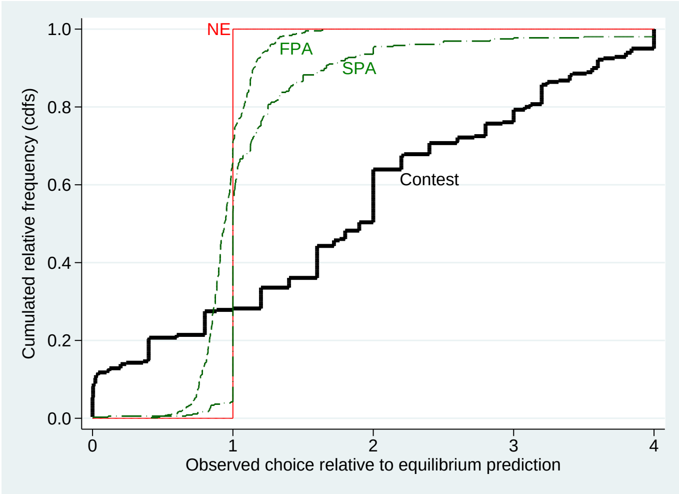

Figure 1: Cumulative Distribution of choices relative to Nash prediction in the Contest (Data from Abbink et al. (2010)), the first-price auction (Data from Brosig and Reiss (2007)), the second-price auction (Data from Cooper and Fang (2008)) and according to Nash equilibrium.

(Fallucchi et al. (2013)). However none of these studies explicitly addresses or explains the large observed behavioral variation. It has been conjectured, though, that heterogeneity in preferences or demographics might be one explanation (Sheremeta (2013)).

In this paper we try to explain both, the large amount of behavioural variation observed and the low explanatory power of Nash equilibrium in these games. We distinguish between two classes of alternative explanations: one based on preferences, the other on bounded rationality. Potentially, there are two sources of complexity that can inhibit the explanatory power of NE if agents are boundedly rational: (i) the difficulty to form correct beliefs and (ii) the difficulty to formulate best responses (even if beliefs were correct). In our experiment we vary these two sources of complexity systematically. If complexity is indeed the underlying reason for the large behavioural variation and the low explanatory power of NE in this game, then we should see (approximate) Nash behavior once both these sources of complexity are removed.[2]

In our benchmark treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$ participants play a standard Tullock contest. Treatment $\mathbf{E}_{asy}\mathbf{D}_{iff}$ exogenously manipulates the difficulty to formulate best responses. This treatment coincides with $\mathbf{D}_{iff}\mathbf{D}_{iff}$ except for the fact that prize allocation is deterministic. In the light of ample evidence that people find it difficult to reduce uncertainty (Kahneman and Tversky (1972)), this should make it easier to formulate best responses. Treatment $\mathbf{D}_{iff}\mathbf{T}_{riv}$ exogenously manipulates the difficulty of forming correct beliefs. It coincides with $\mathbf{D}_{iff}\mathbf{D}_{iff}$, but participants in this treatment play against computer opponents that play pre-determined actions that are announced to participants before they make their choices. Hence there is no strategic uncertainty in this treatment and forming beliefs about the opponent's choices is trivial. Finally

[2]Since we do not observe preferences directly, it is hard to say what the "right" Nash equilibrium benchmark is. Virtually all of the existing literature uses the risk neutral NE as a benchmark (see survey below). We also use risk neutral expected utility maximizers as a benchmark, but, in addition, we also consider a wide range of other preferences. The classes of preferences we study include different risk preferences and social preferences.

treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$ coincides with $\mathbf{D}_{iff}\mathbf{T}_{riv}$, but prize allocation is again deterministic. Hence in this treatment both best response formulation and belief formation should be easy or trivial, respectively.

Our main findings are as follows. As in previous experiments, investments in the standard contest ($\mathbf{D}_{iff}\mathbf{D}_{iff}$) are spread across the whole strategy space and Nash equilibrium has almost no explanatory power. When only one source of difficulty is removed ($\mathbf{E}_{asy}\mathbf{D}_{iff}$ or $\mathbf{D}_{iff}\mathbf{T}_{riv}$), choices are still very different from the Nash prediction, more than half of the time players choose strategies that are strictly dominated and investments choices are highly variable. However, when both best response and belief formation are easy or trivial, respectively ($\mathbf{E}_{asy}\mathbf{T}_{riv}$), players choose strategies that are very close to the prediction of the risk neutral model and behavioural variation is much lower. Bounded rationality can explain the large behavioural variation and low explanatory power of Nash equilibrium in this game.

In order to understand possible interaction effects between our treatment manipulations and participants' (social) preferences, we also conducted a treatment, where we replace computers by humans in treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$ and show that our treatment rankings are unaffected by this change. We also show that all our treatment comparisons and conclusions remain valid if we allow for risk-aversion, risk-seeking preferences or different types of social preferences.

We then ask whether heterogeneity in preferences can explain the large behavioural variation in the standard contest. To these ends we conduct extensive simulations where we consider many possible population compositions of agents with differing risk, social, joy of winning and standard preferences and show that none of these is able to recover the behavioural variation found in the standard contest ($\mathbf{D}_{iff}\mathbf{D}_{iff}$). We conclude that, while some preference based explanations can explain the so-called "overbidding phenomenon", they cannot explain the large behavioural variation typically found in these games. Reducing the complexity of the environment, on the other hand, eliminates the large behavioural variation and leads to behaviour that is very consistent with the predictions of the risk neutral model.

Complexity has been found to play an important role in other games and decision-problems. Grimm and Mengel (2012) show that players are able to learn the Nash equilibrium in normal form games with a unique pure strategy Nash equilibrium in situations of low complexity (few games and easy access to feedback), but not in situations with higher complexity (many games or difficult access to feedback.) Huck et al. (2010) show that participants have more difficulty to form correct beliefs as the environment gets more complex. In decision problems Huck and Weizsäcker (1999) find that players are more likely to deviate from expected value maximization when choosing between a pair of lotteries if the task is more complex, as measured by the number of possible outcomes (see also Rabin and Weizsäcker (2009)).

The paper is organized as follows. In section 2 we present the experimental design. Section 3 lists the conjectures that are tested in section 4, where we show how reducing complexity reduces behavioural variation and increases the explanatory power of Nash equilibrium. Preference based explanations and heterogeneity in preferences are discussed in section 5 and conclusions are drawn in section 6. An Appendix contains additional table and figures as well as the experimental instructions.

2 Experimental Design

Our experimental design aims to understand whether bounded rationality could be the reason for the large behavioural variation and low explanatory power of Nash equilibrium in Tullock contests. In order to reach a Nash equilibrium, players need to do two things: (i) they need to be able to form correct beliefs and (ii) they need to be able to formulate best responses (to whatever beliefs they hold). Participants could learn to form correct beliefs or formulate best responses in a number of ways. Under some learning models they might even learn to play a Nash equilibrium without explicitly forming beliefs at all (e.g. reinforcement learning). How exactly participants do so is of secondary concern for us in this study. The question we ask is whether if we make it "easy enough" to formulate beliefs and best responses behavioural variation in this game will be (substantially) reduced and the explanatory power of Nash equilibrium increased. In a 2×2 factorial design we hence varied complexity along these two dimensions. Table 1 summarizes this treatment structure.

Treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$ is our benchmark treatment and implements the standard contest as typically studied in the experimental literature. In treatment $\mathbf{E}_{asy}\mathbf{D}_{iff}$, we make it "easy" to formulate best responses, but keep the difficulty of forming correct beliefs. This reverses in treatment $\mathbf{D}_{iff}\mathbf{T}_{riv}$, where we keep the difficulty of formulating best responses and make it simple, in fact trivial, to formulate correct beliefs. In treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$, both sources of complexity are eliminated. Next, we describe these treatments in detail.

	Belief Formation	
Best Response Formulation	Difficult	Trivial
Difficult	$\mathbf{D}_{iff}\mathbf{D}_{iff}$(54)	$\mathbf{D}_{iff}\mathbf{T}_{riv}$(44)
Easy	$\mathbf{E}_{asy}\mathbf{D}_{iff}$(54)	$\mathbf{E}_{asy}\mathbf{T}_{riv}$(44); $\mathbf{E}_{asy}\mathbf{T}_{riv}{-}h$ (48)

Table 1: Experimental Design, numbers in brackets indicate the total number of participants in each treatment. Participants in treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$ were allocated to matching groups of 6 participants, thus there are 9 independent observations in each of these treatments. In treatments $\mathbf{E}_{asy}\mathbf{T}_{riv}$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}$ each participant is an independent observation. In $\mathbf{E}_{asy}\mathbf{T}_{riv}{-}h$ 24 participants were "dummy" players, so we use the choices of the remaining 24 players in our analysis.

$\mathbf{D}_{iff}\mathbf{D}_{iff}$ In treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$, subjects participated in the standard Tullock contest. In our Tullock contest, two players compete for a commonly known prize of 16 Experimental Currency Units (ECU). In every round, participants received an endowment of 16 ECU. Participants could then invest an amount from this endowment, i.e. an amount from the action set $A_i = \{1, 2, ..., 16\}$ with typical element a_i $(i = 1, 2)$. Players' investments are sunk costs and the sets of feasible monetary payoffs are given by $\Pi_i = \{16 - a_i, 16 + 16 - a_i\}$. The contest success function that denotes the probability that player i receives the prize and, hence, the payoff of $(32 - a_i)$ ECU is given by $\rho_i(a_1, a_2) = \frac{a_i}{a_1+a_2}$. The experimental instructions can be found in Appendix B.

The unique (risk-neutral) Nash equilibrium in our contest game is given by $(a_1^*, a_2^*) = (4, 4)$. This is the unique strategy profile surviving iterated elimination of dominated strategies. All choices $a_i > 4$ are strictly dominated by $a_i = 4$ under the standard assumptions. Put differently, there are no beliefs that can support choices above the Nash choice of 4 for risk neutral expected utility maximizers. Even allowing for moderate degrees of risk aversion or risk seeking does not affect this property. To demonstrate this, Figure 6 in the Appendix illustrates the best response correspondence for a risk averse, risk neutral and risk-seeking CRRA agent. These best response functions are very similar to one another and choices exceeding 4 are not rationalizable for any of these types. Therefore, we will follow the vast majority of the literature and use the risk-neutral Nash equilibrium as a benchmark. We will be very careful, though, to account for heterogeneity in risk and social preferences when interpreting our results and we will investigate this issue more deeply in section 5.

$\mathbf{E}_{asy}\mathbf{D}_{iff}$ In treatment $\mathbf{E}_{asy}\mathbf{D}_{iff}$ we wanted to make it easier for participants to formulate best responses (compared to $\mathbf{D}_{iff}\mathbf{D}_{iff}$). An obvious treatment manipulation would be to simply compute the best response for them. Remember, though, that we want to allow for the possibility that participants may have heterogeneous preferences. If we want to remain agnostic about their preferences, we cannot compute best responses "for them". In the light of ample evidence that people have difficulty to reduce uncertainty (e.g. Kahneman and Tversky (1972); Peters et al. (2007)), one treatment variation that should make it easier to formulate best responses is to eliminate random variables.[3] One way of doing so is to share the prize between both players according to their individual investment shares in total investment. Hence, instead of receiving $(32 - a_i)$ ECU with probability $\frac{a_i}{a_1+a_2}$ and $(16 - a_i)$ ECU with complementary probability (as in $\mathbf{D}_{iff}\mathbf{D}_{iff}$), subjects received as the monetary round payment the hypothetical expected value

$$\frac{a_i}{a_1 + a_2}(32 - a_i) + (1 - \frac{a_i}{a_1 + a_2})(16 - a_i) = 16 - a_i + \frac{16a_i}{a_1 + a_2} \text{ ECU.}$$

In this treatment - given correct beliefs - formulating best responses should be easier, because participants do not have to reduce uncertainty. It is still difficult in this treatment, though, to form correct beliefs.

$\mathbf{D}_{iff}\mathbf{T}_{riv}$ Making it simpler (or even trivial) to form correct beliefs requires more elaborate design interventions. If one would like to maintain the simultaneous choice setting (which is desirable because of treatment comparisons), the only way is to let participants play against computers or against human players who are so restricted in their choice of strategy that they could almost be replaced by computers. In the experiment we did both. To be able to eliminate strategic uncertainty in treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$, all treatments (including the benchmark treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$) were divided into 4 successive blocks and each block was composed of 10 rounds of play.

[3] In section 5.1, we analyze the possibility that this treatment variation interacts with risk preferences to generate differential results.

Specifically, strategic uncertainty was removed by matching participants to computers who played a fixed strategy in every block.[4] Players were informed about the opponent's choices ex ante and, hence, should hold deterministic and, moreover, correct beliefs about the opponent's play. The strategy adopted by the computer was held constant within each block of ten rounds and changed from one block to the other, thus, each player faced four different strategies. The four strategies used by the computer were randomized across participants who faced one of the following sequences where each computer action was played in ten rounds: (1, 14, 11, 8), (5, 10, 3, 16), (9, 6, 15, 4) or (13, 2, 7, 12). Each sequence was allocated to the same number of players. Hence, there is an equal number of observations for each computer choice in $\{1, ..., 16\}$. Sequences were selected such that each participant faced the same average level of computer investment and some "high" as well as "low" investment levels.

$\mathbf{E}_{asy}\mathbf{T}_{riv}$ In treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$ both sources of complexity were eliminated using the same procedures as adopted for treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$and $\mathbf{E}_{asy}\mathbf{D}_{iff}$, respectively. Since players were ex ante informed about the investment choice of the computer, they should hold correct beliefs in $\mathbf{E}_{asy}\mathbf{T}_{riv}$(just as in $\mathbf{D}_{iff}\mathbf{T}_{riv}$). The question is whether - given correct beliefs - participants engage in best response behaviour. This should be easiest in $\mathbf{E}_{asy}\mathbf{T}_{riv}$where *both* sources of complexity are eliminated.

$\mathbf{E}_{asy}\mathbf{T}_{riv}$-h Treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$-***h*** coincides with treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$ except that computers were replaced by human players that were restricted to choose the same strategy in each block of ten rounds where the strategies. Moreover, subjects had to choose from a set of four different strategies. The four different sets were $\{1, 14, 11, 8\}, \{5, 10, 3, 16\}, \{9, 6, 15, 4\} or \{13, 2, 7, 12\}$. To give the human players some choices in treatments $\mathbf{E}_{asy}\mathbf{T}_{riv}$-*h*, they could pick the order in which they played the strategies, while in $\mathbf{E}_{asy}\mathbf{T}_{riv}$the computer played them in a fixed sequence not known to the human participants. Essentially, treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$-*h* serves as a robustness check to understand whether introducing computers instead of human subjects affects our treatment comparison, e.g., due to social preferences. This is, by far, not the only robustness check we do, however to account for this possibility (see section 5.2).

Other Details To give subjects the opportunity of experimenting and learning about the game situation at no cost, during the experiment players were facing an alternating pattern of 5 non-incentivized and 5 incentivized rounds. This design was chosen to encourage participants to practice and understand the design and incentives well. After all 40 rounds were completed, players took an incentivized numeracy test that measured the ability to understand and manipulate probabilities.[5] At the end of the experiment, one incentivized round was randomly chosen

[4]A similar procedure was used e.g. by Johnson et al. (2002), who match players to "robots" in three-round bargaining game. Robots were playing according to the subgame-perfect Nash equilibrium prediction and induced human players to choose strategies that are closer to the theoretical prediction.

[5]Questions for the numeracy test were taken from Peters et al. (2007). From the 15 questions in the original test we removed 7 questions that were found too simple (correctly answered by at least 80% of the population with higher education). We did so to increase incentives of answering the more complicated questions and to

from each block and participants received the earnings from the chosen rounds as well as the payment for correct answers in the questionnaire. The experiment was conducted in March and September 2012 using z-Tree (Fischbacher, 2007) and ORSEE (Greiner, 2004) at the BEElab at Maastricht University. A total of 244 students participated in the experiment. The average duration of the experiment was 60 minutes and participants on average earned 15.15 euros.

3 Conjectures

Our first Conjecture is that reducing complexity by reducing the difficulty of formulating best responses and forming correct beliefs leads to less behavioural variation.

Conjecture 1. *Behavioural Variation is highest in* $\boldsymbol{D}_{iff}\boldsymbol{D}_{iff}$ *and lowest in* $\boldsymbol{E}_{asy}\boldsymbol{T}_{riv}$.

We introduce our measure of behavioural variation further below (Section 4). Before we do that, we formulate conjectures about the explanatory power of NE in our four treatments. Since we induce non-Nash choices by computerized players in the treatments with trivial belief formation, $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$, the percentage of choices which are Nash is not a reasonable measure. Instead we decompose Nash behaviour into the two components that motivated our design: (i) correct beliefs and (ii) best response behaviour. In the following we explain how our treatment variations impact each of these components.

We start by looking at the effect of making it easier to formulate best response. There is a large body of evidence that people have difficulty in comprehending probabilistic statements and in making decisions in probabilistic environments (Kahneman and Tversky (1972)). If NE has explanatory power, then we would expect that removing this uncertainty facilitates best response behaviour. As a consequence, we expect a larger share of choices to be consistent with best response behaviour to *some* beliefs when best response formulation is easy, i.e. in $\mathbf{E}_{asy}\mathbf{D}_{iff}$ compared to $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ compared to $\mathbf{D}_{iff}\mathbf{T}_{riv}$. What does it mean to be "consistent with best response behaviour"? It means that choices are rationalizable, i.e. that there *exists* a belief that would justify the player's choice. Since dominated choices are *not* consistent with best response behaviour, we expect to see less of these choices when best response formulation is easy, as in $\mathbf{E}_{asy}\mathbf{D}_{iff}$ compared to $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ compared to $\mathbf{D}_{iff}\mathbf{T}_{riv}$. Of course, what constitutes a dominated strategy depends on the agent's preferences. We follow the vast majority of the literature and assume a risk neutral expected utility maximizer as a benchmark, but we carefully look at risk and social preferences in Section 5. There, we demonstrate that our results obtain for very large classes of preferences.

Conjecture 2 *A larger share of choices is undominated with easier best response formulation, i.e. in* $\boldsymbol{E}_{asy}\boldsymbol{D}_{iff}$ *compared to* $\boldsymbol{D}_{iff}\boldsymbol{D}_{iff}$ *and in* $\boldsymbol{E}_{asy}\boldsymbol{T}_{riv}$ *compared to* $\boldsymbol{D}_{iff}\boldsymbol{T}_{riv}$.

Furthermore, if the difficulty in dealing with probabilistic statements is the main reason participants make suboptimal decisions in non-deterministic environments, then participants that are

enable better differentiation based on performance. Subject were paid 1 ECU (0.25 euro) for a correct answer for each of the first 7 questions, and 2 ECU (=0.5 euro) for a correct answer to the final question that was the most difficult one.

better at understanding probabilistic environments should do "better". In other words subjects with higher scores in the numeracy test should choose strategies consistent with best response behaviour more often than others (in treatments $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$). We address this issue in Section 4.3, where we present evidence from our questionnaire.

Since belief formation is trivial in treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$, participants should have "correct" beliefs in these treatments. As a consequence, if NE has explanatory power, then we expect that a larger share of choices to be a best response to the opponent's behaviour with trivial belief formation, i.e. in $\mathbf{D}_{iff}\mathbf{T}_{riv}$ compared to $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ compared to $\mathbf{E}_{asy}\mathbf{D}_{iff}$.

Conjecture 3 *A larger share of choices is a best response to the opponent's behaviour with trivial belief formation, i.e. in $\boldsymbol{D}_{iff}\boldsymbol{T}_{riv}$ compared to $\boldsymbol{D}_{iff}\boldsymbol{D}_{iff}$ and in $\boldsymbol{E}_{asy}\boldsymbol{T}_{riv}$ compared to $\boldsymbol{E}_{asy}\boldsymbol{D}_{iff}$.*

4 Results: Reducing Complexity

In this section, we present our main results. We start by evaluating Conjecture 1 on behavioural variation and compare it across treatments (section 4.1). Subsequently, we focus on the explanatory power of Nash equilibrium regarding its two components, best response behaviour and correct beliefs, and evaluate Conjectures 2 and 3 (section 4.2). Unless explicitly stated otherwise, throughout the paper we analyze data generated in the incentivized rounds during the second half of the experiment. The reason is that we want to focus on mature behaviour and eliminate behavioural variation that would disappear after some learning has occurred.

4.1 Behavioural Variation

Figure 2 is the analogue of Figure 1 and compares treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ in terms of behavioural variation. The figure plots the empirical cumulative distribution functions (cdf) of choices a_i divided by the equilibrium prediction along with the theoretically predicted cdf. Theoretically, hence, all the observations should yield a degenerate cdf where all the cdf mass cumulates at one, because choices should equal the equilibrium prediction in each period. What is the equilibrium prediction for both treatments? As discussed before, in treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$ (i) the risk-neutral Nash equilibrium prescribes an investment of $a_i^* = 4$ ECU and (ii) the best response of a risk-neutral agent to the empirical distribution of choices in $\mathbf{D}_{iff}\mathbf{D}_{iff}$ is also $a_i^* = 4$. In treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$, however, the theoretical prediction is different. Since the predetermined computer's choice can differ from 4 ECU and is announced, hence, known before making a choice, the prediction is simply the best response to that predetermined choice. The resulting best responses range from 1 to 4 depending on the computer's choice, noting that all choices exceeding 4 are strictly dominated.

Figure 2 shows that our benchmark treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$ produces results for the Tullock contest that are "standard" in the sense of high behavioural variation and choices completely disconnected from the theoretical (Nash) prediction. There does not seem to be more cdf mass at 1,

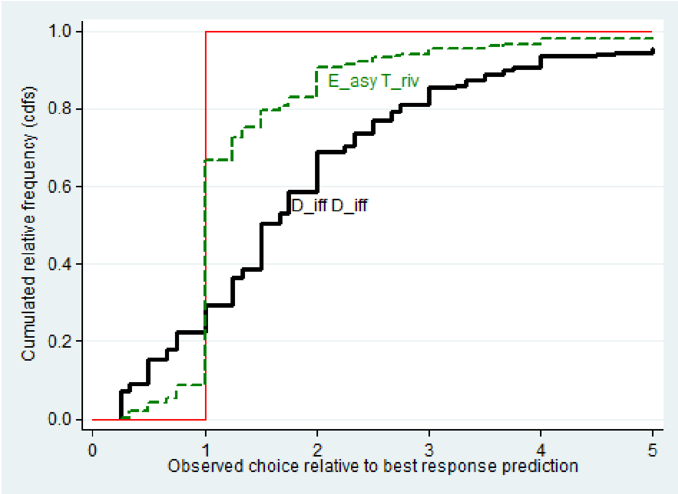

Figure 2: Cumulative distributions of ratio variable, where investment choices are divided by equilibrium prediction, in theory and treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ as well as $\mathbf{E}_{asy}\mathbf{T}_{riv}$.

where the actual choice equals the theory prediction, than elsewhere. Note also that all the cdf mass to the right of 1 stems from dominated choices for a risk-neutral and moderately risk-averse (see section 5.1) agents. In treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$, on the other hand, the theoretical prediction clearly has something to say about the data with most cdf mass concentrated at 1.

We use two measures previously used in the literature to obtain more conventional measures of behavioural variation. The first measure is the standard deviation of choices. The second measure is entropy, which evaluates the stochastic variation of a random variable that can assume a finite set of values (Shannon, 1948). Entropy has been used, e.g. by Bednar et al. (2011), to evaluate behavioural variation in normal-form games. While standard deviation is a very common measure, the advantage of using entropy is that it yields a measure of the average unpredictability of a random variable that captures the amount of information needed to describe a distribution. This is why it captures nicely the difficulty to form correct beliefs. If a is the random variable of investments and $p_i = P(a = a_i)$ is its probability density function for all possible strategies $a_i = 1, 2, \ldots, 16$, entropy is computed as:

$$H = - \sum_{i=1,2,\ldots,16} p_i \log_2(p_i)$$

It is common to use a logarithm to base 2, so that entropy can be interpreted as the total number of bits needed to describe the data. For our strategy space with 16 possible choices, the entropy measure can take values from 0 (if a single strategy is always chosen) to 4 (if all strategies are chosen with equal frequency).[6]

To compare behavioural variation across treatments, we compute our measures (entropy and standard deviation) by treatment conditional on the strategy chosen by the opponent. Aggregate entropy in each treatment is then computed as the weighted average of the conditional entropy levels and the weights are determined by the frequency with which each of these strategies

[6]To accommodate the case that a strategy is never played, we follow the literature and assume $0 \cdot \log(0) = 0$.

are chosen in the standard contest ($\mathbf{D}_{iff}\mathbf{D}_{iff}$). Conditioning on opponents' strategies ensures that the behavioural variation is zero in all treatments if all participants always choose best responses to correct beliefs. Without this conditioning, there would be artificial variation in treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$, which is merely due to the computer changing strategies between blocks, which in some cases changes the best response.[7]

	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$
Entropy	3.21	2.79	2.45	1.50
Std. dev.	3.27	2.42	3.15	1.16

Table 2: Behavioral variation across treatments. Data from incentivized rounds in the the second half of the experiment.

Both measures are summarized in Table 2. Both entropy and standard deviation are highest in the standard contest ($\mathbf{D}_{iff}\mathbf{D}_{iff}$) and lowest when complexity is lowest ($\mathbf{E}_{asy}\mathbf{T}_{riv}$). The two measures disagree on ranking the intermediate treatments, where either best response formulation or belief formation is simplified (but not both): entropy is higher in $\mathbf{E}_{asy}\mathbf{D}_{iff}$ than in $\mathbf{D}_{iff}\mathbf{T}_{riv}$, but standard deviation ranks them the other way round. Table 12 in Appendix A.6 shows that both within and between subject variability contribute to the high amount of behavioural variability observed in $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and that eliminating uncertainty (as in $\mathbf{E}_{asy}\mathbf{T}_{riv}$) reduces both types of variability.

We conduct Wilcoxon ranksum (Mann-Whitney) tests to check for the statistical significance of these differences. To account for the dependency of observations within matching groups we compute entropy (standard deviation) separately for each matching group and then compare the distribution of the results using two-sided ranksum tests. Since in treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ each participant is an independent observation, we can either compare matching group averages with individual levels or form artificial matching groups in the $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$ treatments and then compare matching groups. We do both and report the lower/higher p-value, whichever is more relevant. For entropy, we find that the difference between $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ is significant ($p < 0.0544$). Hence, reducing both dimensions of complexity significantly reduces behavioural variation. If best response formulation is easy, making belief formation trivial also significantly reduces behavioural variation. The difference between $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ is significant ($p < 0.0210$). The comparisons between the intermediate treatments $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}$ are never significant ($p > 0.7003$), so that both one-dimensional reductions of complexity affect behavioural variation similarly.

In terms of the standard deviation we again find significant differences between $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}(p < 0.0005)$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}(p < 0.0047)$ and no statistical significance when comparing $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}(p > 0.2110)$. Overall we find clear support for Conjecture 1 and conclude that reducing complexity along the two dimensions in our design does reduce behavioural variation.

[7] Despite this fact - even if entropy and standard deviation are computed by simply pooling all the observations in a treatment, treatments are ranked similarly. The only exception would be treatment $\mathbf{D}_{iff}\mathbf{T}_{riv}$ which in that case exhibits larger behavioural variation than $\mathbf{D}_{iff}\mathbf{D}_{iff}$.

Result 1 *Behavioral variation is lowest in* $\boldsymbol{E}_{asy}\boldsymbol{T}_{riv}$ *, where best response formulation is easy and belief formation trivial, among all treatments and it is lower in* $\boldsymbol{E}_{asy}\boldsymbol{D}_{iff}$ *and* $\boldsymbol{D}_{iff}\boldsymbol{T}_{riv}$ *, where one dimension of complexity is reduced, compared to* $\boldsymbol{D}_{iff}\boldsymbol{D}_{iff}$.

4.2 Explanatory Power of Nash Equilibrium

Now that we have established that making best response formulation easy and belief formation trivial reduces behavioral variation substantially, we examine whether the explanatory power of Nash equilibrium responds accordingly. We evaluate this conjecture in detail in this section. We follow the vast majority of the literature and focus on the benchmark prediction for risk neutral agents. In Sections 5.1 and 5.2 we relax this assumption and explore different preference specifications, in particular different risk preferences and social preferences.

	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$
$P(a = NE)$	7.04%	13.33%	-	-
$P(a = BR)$	7.04%	12.04%	22.50%	65.23%
$P(\|a - NE\| \leq 1)$	25.74%	32.78%	-	-
$P(\|a - BR\| \leq 1)$	26.30%	31.30%	47.95%	83.64%
$P(a > 4)$	60.19%	62.78%	51.36%	16.14%

Table 3: Indicators measuring the explanatory power of Nash equilibrium across treatments. Data from the incentivized rounds in the second half of the experiment.

Table 3 shows summary statistics on different measures of the explanatory power of Nash equilibrium. Let us first focus on the benchmark treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$. In this treatment, we observe only 7 percent of choices equal to the Nash prediction. To get a sense of how little this is, note that a player choosing uniformly at random between the possible investment levels $a_i \in \{1, ..., 16\}$ would play the Nash choice 6.25 percent of the time. The explanatory power of NE is not much improved when we look at the percentage of choices that are "close" to the Nash choices ($a_i \in \{3, 4, 5\}$). This percentage is barely above 25 percent. Note again that a random player would hit these numbers in 18.75 percent of the cases. Moreover, approximately 60 percent of all choices in $\mathbf{D}_{iff}\mathbf{D}_{iff}$ are strictly dominated ($a > 4$). The top-left panel of Figure 3 illustrates these results, where the dashed line separates dominated ($a > 4$) from undominated choices ($a \leq 4$). Our results in the benchmark treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$ hence display the low explanatory power of NE typically found in Tullock contest data (Abbink et al. (2010), Sheremeta (2013)). We now study how it changes when we reduce the difficulty of (i) best response formulation and (ii) belief formation.

"Easy best responses" (Conjecture 2) Let us first see whether and how the explanatory power of NE is increased by making it easier to formulate best responses. As outlined in Section 3, we expect participants to choose strategies that are consistent with best responses to *any belief*, i.e. undominated strategies, more often when best response formulation is easier, i.e. in $\mathbf{E}_{asy}\mathbf{D}_{iff}$ compared to $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ compared to $\mathbf{D}_{iff}\mathbf{T}_{riv}$. For risk-neutral (and moderately risk averse) agents, investment levels below or equal 4 are undominated. Table

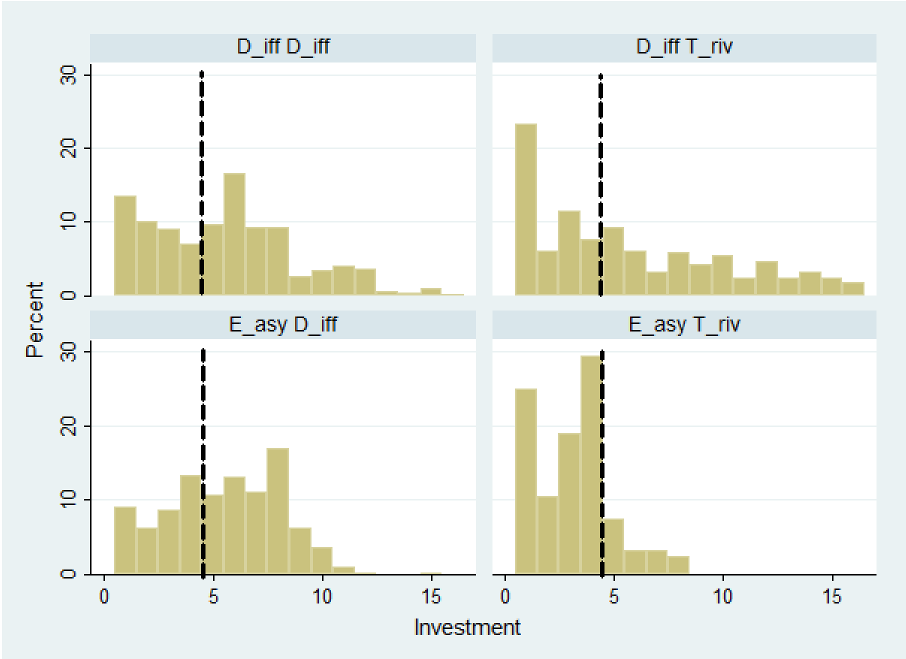

Figure 3: Distribution of investment choices by treatment in the incentivized rounds of the second half of the experiment (rounds 26-30 and 36-40). The vertical dashed lines separate undominated (≤ 4) from dominated choices (> 4).

3 shows the percentage of dominated choices by treatment. If belief formation is difficult, then making best response formulation easy does not improve the explanatory power of NE, as the comparison of treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$ shows. In fact the percentage of strictly dominated strategies is even slightly higher under $\mathbf{E}_{asy}\mathbf{D}_{iff}$ compared to $\mathbf{D}_{iff}\mathbf{D}_{iff}$, though this difference is not statistically significant (see Table 8 in Appendix A.2). Overall, the explanatory power of Nash equilibrium is very low in both treatments. This is also illustrated by Figure 3 (top-left vs bottom-left panel).

If belief formation is trivial, on the other hand, we observe a strong effect, as shown in Figure 3 (top-right vs bottom-right panel). Comparing treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$, we find that making best response formulation easy, substantially improves NE's explanatory power. While 51.36% of choices are strictly dominated in $\mathbf{D}_{iff}\mathbf{T}_{riv}$, only 16.14% are so in $\mathbf{E}_{asy}\mathbf{T}_{riv}$(see Table 3). This difference is highly statistically significant ($p < 0.0001$) according to a simple logit regression reported in Table 8 in Appendix A.2. Hence making it easier for participants to formulate best responses, reduces the frequency of dominated choices if and only if it is trivial to form correct beliefs.

Figure 4 illustrates this effect in more detail. The figure plots the quartiles of responses to each of the strategies chosen by the computer. The difference between the first and the third quartiles is much larger in treatment $\mathbf{D}_{iff}\mathbf{T}_{riv}$, especially for large investment values chosen by the computer. In treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$ median investments track the best response function well.[8]. Table 9 in Appendix A shows that differences between treatments do not disappear

[8]Additional evidence for this effect can be found in Table 10 in Appendix A that compares the median

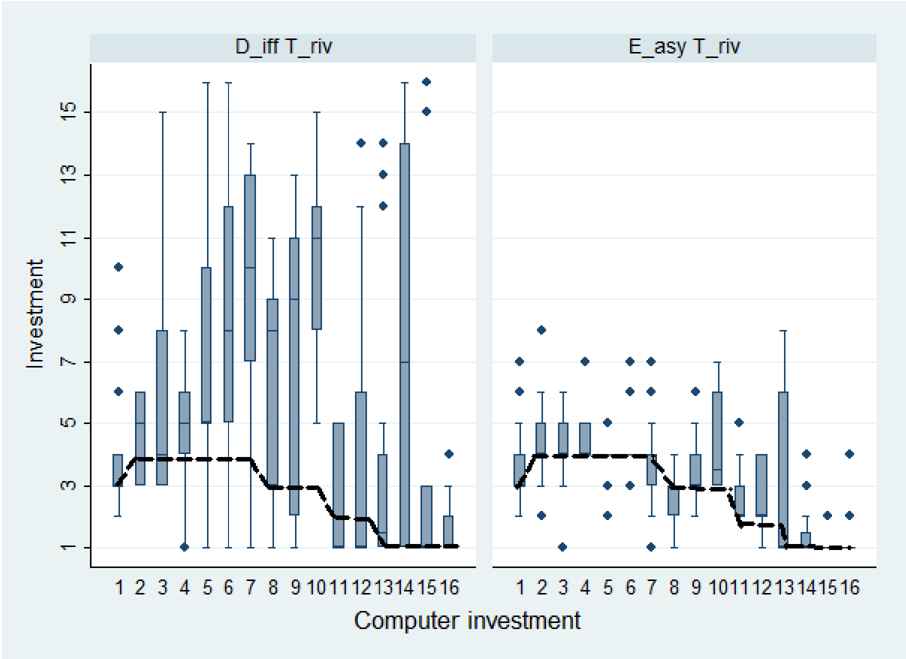

Figure 4: Best-response (dashed line) and the empirical distribution of investments, conditional on the strategy played by the computer. Box plots represent the medians, differences between 25th and 75th percentiles and lower and upper adjacent values.

with learning. The difference between the $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$ is insignificant in every block except for the first one where it is marginally significant, while the difference between $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ is highly significant throughout the experiment.

To sum up, there is a complementary effect. Making it easier to formulate best responses reduces the frequency of dominated choices if and only if it is trivial to form correct beliefs.

Result 2 *Making best response formulation easy, increases the explanatory power of NE if and only if belief formation is trivial. More choices are undominated under $\boldsymbol{E}_{asy}\boldsymbol{T}_{riv}$ compared to $\boldsymbol{D}_{iff}\boldsymbol{T}_{riv}$.*

"Trivial Beliefs" (Conjecture 3) If beliefs are correct, then the explanatory power of Nash equilibrium is reflected in the share of choices which are best responses to the opponent's behaviour. Recall, that in the treatments with trivial belief formation the opponent's behaviour is announced, giving rise to correct beliefs. For comparison to the treatments with difficult belief formation we report the frequency of choices that are best response to actual opponents' behaviour in the latter.[9]

Consistently with Conjecture 3, Table 3 shows that the share of best responses increases as belief formation becomes trivial irrespective of whether best response formulation is difficult or easy.

investments for each choice of the opponent across treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$. It can be seen that the median investment level tends to be closer to the best response in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ and that the treatment difference is particularly large for situations where the opponent chooses a very high investment level (5 and above)

[9]Results are qualitatively the same if we consider best responses to population averages, instead.

While less than 10% of participants best respond to their opponent's choices in $\mathbf{D}_{iff}\mathbf{D}_{iff}$ this number more than doubles in $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and reaches 22.50%. Table 8 in Appendix A.2 shows that the treatment difference is significant ($p < 0.0001$). The effect is even stronger if we compare $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$. Here the share of choices which are best responses to the opponent's behaviour increases from about 12% to more than 65%. Also this treatment difference is statistically significant ($p < 0.0001$, Table 8). 83.64% of choices in $\mathbf{E}_{asy}\mathbf{T}_{riv}$deviate from the best response by 1 ECU or less.[10]

Result 3 *Making belief formation trivial, increases the explanatory power of NE. More choices are best responses to the opponent's behaviour under $\boldsymbol{E}_{asy}\boldsymbol{T}_{riv}$ compared to $\boldsymbol{E}_{asy}\boldsymbol{D}_{iff}$ and under $\boldsymbol{D}_{iff}\boldsymbol{T}_{riv}$ compared to $\boldsymbol{D}_{iff}\boldsymbol{D}_{iff}$.*

It should be noted that the difference between $\mathbf{E}_{asy}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$ is larger than the difference between $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{D}_{iff}\mathbf{D}_{iff}$, i.e. making belief formation trivial is more effective in increasing the explanatory power of NE if best response formation is easy ($\chi^2 : p < 0.0023$, see Appendix A.2). Recall, also that we have seen that making best response formation easy is effective only if belief formation is trivial (no significant difference between $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{D}_{iff}\mathbf{D}_{iff}$). Hence both sources of complexity compound to decrease the explanatory power of NE. Removing both sources of complexity leads to very high consistency with best response behavior and a consequently high explanatory power of Nash equilibrium. These results suggest that bounded rationality is a key force behind the typically low explanatory power of NE in Tullock contests.

4.3 Cognitive Abilities and Equilibrium Deviations

In the previous subsection we have established that, while behavioural variation is huge in $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and significantly reduced in $\mathbf{E}_{asy}\mathbf{T}_{riv}$, there remains a substantial amount of behavioural variation in $\mathbf{D}_{iff}\mathbf{T}_{riv}$. This shows that, even if belief formation is trivial, participants still find it difficult to formulate best responses with probabilistic prize allocation. Since in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ (where prize allocation is deterministic) this variation disappears, the difficulty in forming best responses seems to be related to (some) participants' difficulty in dealing with probabilistic choice situations.

In this subsection we evaluate this hypothesis using data from our post-play questionnaire on cognitive ability and risk numeracy.[11] If complexity is indeed one of the main reasons why participants' choices are not consistent with best response behaviour in the experiment, then one might conjecture that those with better scores in a risk numeracy test should have less difficulty in dealing with this complexity and hence should be closer to best response behaviour.

As argued above, we expect risk numeracy to matter particularly in treatment $\mathbf{D}_{iff}\mathbf{T}_{riv}$, where belief formation is trivial, but best response formation difficult. It is not clear why cognitive

[10] $\mathbf{E}_{asy}\mathbf{T}_{riv}$ is the treatment where most learning is observed (see Table 9 in Appendix A). The change in average equilibrium deviations decreases by 63% from the first to the fourth block.

[11] The entire questionnaire, as well as summary statistics on answers can be found in Appendix D.

ability (as measured in numeracy tests) should help to form correct beliefs, i.e. explain variation in $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$. In treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$, on the other hand, where both best response formation and belief formation are easy we have seen that behavioural variation is low. We focus hence on $\mathbf{D}_{iff}\mathbf{T}_{riv}$. Figure 5 plots the score in the numeracy test (higher score means higher numeracy) against the distance a participant's choice has on average from the best response (Figure 5(a)) and against the computer strategy (Figure 5(b)).

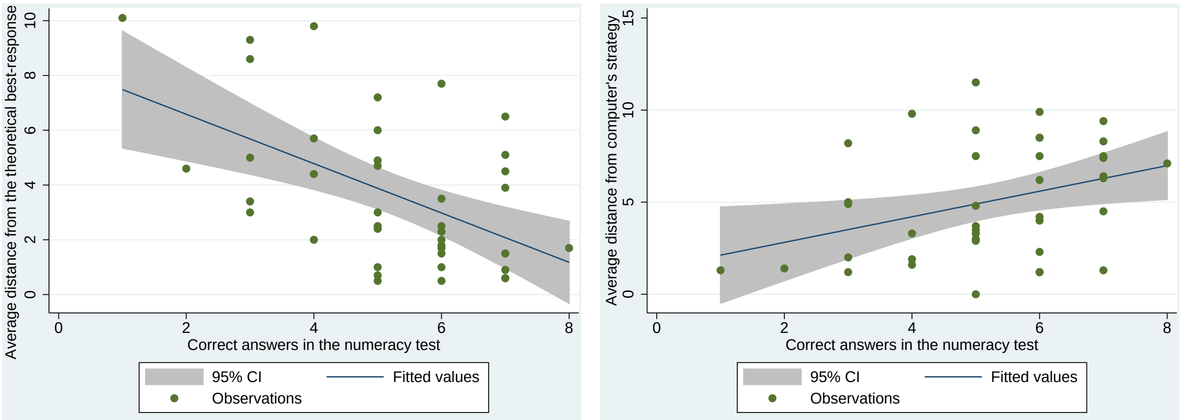

(a) Distance to BR

(b) Distance to Computer Strategy

Figure 5: Relationship between the performance in a numeracy test and absolute difference between investments and computer's strategy/best response, averaged by subjects. Incentivized rounds from the second half of the experiment.

The figure shows data points as well as predicted values from an OLS regression and the associated 95 percent confidence interval. It can be seen quite clearly that participants with higher scores in the numeracy test are closer on average to best response behaviour. They are also less prone to imitating the computer strategy, as Figure 5(b) suggests, but the statistical significance here is somewhat lower. The underlying regression tables can be found in Appendix D. Both results provide further support to our conjecture that the difficulty of participants to make decisions in uncertain environments is a main cause for the low explanatory power of Nash equilibrium in the Tullock contest.

Result 4 *Participant who do better in the numeracy test best respond more often and imitate less often than those who do worse in treatment* $\boldsymbol{D}_{iff}\boldsymbol{T}_{riv}$.

5 Preference Based Explanations

In this section we first ask whether our treatment comparisons and conclusions remain valid under alternative assumptions on preferences. We first look at risk preferences (Section 5.1) and then at social preferences (Section 5.2). Afterwards, in Section 5.3 we ask whether heterogeneity in preferences can explain the large behavioural variation observed in the standard contest.

5.1 Risk Preferences

In this subsection we show that our results and interpretation are robust to considering different risk preferences. First, note that results on behavioural variation are independent of our participant's preferences. In this section, we hence focus on our results from Section 4.2 (on the explanatory power of Nash equilibrium) and show that they are robust when different risk preferences are considered.

The first question to ask is how risk aversion (or risk seeking behaviour) affects the theoretical predictions in the standard contest. Figure 6 in Appendix A shows the best response function for a risk-seeking (parameter $r = -0.5$), risk-neutral ($r = 0$) and risk-averse ($r = 0.5$) CRRA agent. The figure shows that the best response to choices below or equal to four are almost identical for these three types. It further shows that all choices above 4 are never a best response for any of these types. In addition, the more risk-averse an agent is, the more the best-response function shifts downwards, i.e. the lower his best response is. Hence choices above the Nash level of 4 cannot be explained by risk aversion (see also Hillman and Katz (1984) or Abbink et al. (2010)). Table 4 shows that the percentage of choices consistent with best response behaviour are much higher in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ compared to all other treatments also for risk averse and risk-loving agents and they remain very low (barely above uniformly random) in treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$. Further, risk aversion and in particular also risk-seeking preferences are even less consistent with evidence in $\mathbf{E}_{asy}\mathbf{T}_{riv}$, suggesting that risk aversion did not play a major role in the experiment. It has also been shown that loss aversion or S-shape probability weighting cannot explain such choices Baharad and Nitzan (2008). This also means that the results shown in Table 3 would remain unchanged if moderately risk averse or risk-loving CRRA agents are considered.

Treatment	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$
Risk neutrality	7.04%	12.04%	22.50%	65.23%
Risk aversion (r=0.5)	8.89%	10.56%	25.12%	45.00%
Risk seeking (r=-1)	10.00%	13.06%	8.75%	26.25%

Table 4: Percentage of choices that are best responses to the opponent's choice ($\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$) or the empirical distribution of choices ($\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$) for risk averse (CRRA with $r = 0.5$)/risk-neutral/risk-seeking (CRRA with $r = -1$) players.

A different question is whether, in the presence of risk averse agents, our treatment comparisons are distorted by making prize allocation deterministic, as this reduces risk for our participants. In particular a risk-averse CRRA agent might respond to the risk associated with probabilistic prize allocation in treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}$, but behave as if he was risk neutral in $\mathbf{E}_{asy}\mathbf{T}_{riv}$, because all risk has been eliminated there. The preceding discussion suggests that - given moderate degrees of risk aversion - this should not make too much of a difference. In particular, Figure 6 in Appendix A shows that under CRRA utility risk averse participants invest *more* than risk neutral participants if the opponent invests very little and *less* than risk neutral participants if the opponent invests a lot. Since in treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$, participants know their opponents' investments, we can use this to test for such distortions. If risk aversion plays a significant role in explaining the differences between these treatments, then we should expect that investments are higher in $\mathbf{D}_{iff}\mathbf{T}_{riv}$ compared to $\mathbf{E}_{asy}\mathbf{T}_{riv}$ for low investment choices of the opponent and lower for high investment choices of the opponent. Figure 4 shows that this is not the case. Mostly behavioural variation is higher under $\mathbf{D}_{iff}\mathbf{T}_{riv}$ compared to $\mathbf{E}_{asy}\mathbf{T}_{riv}$, but if at all investments in $\mathbf{D}_{iff}\mathbf{T}_{riv}$ are lower for low investments and higher for high investments of the opponent compared to $\mathbf{E}_{asy}\mathbf{T}_{riv}$. Hence we can rule out this type of distortion of our treatment comparisons.

We conclude that the results summarized in Table 3 as well as treatment rankings are robust to considering moderate degrees of risk aversion or risk seeking behaviour.

5.2 Social Preferences

Similarly we ask whether our conclusions drawn so far, and in particular our treatment comparisons, are valid if our participants had social preferences (inequity aversion, joy of winning, reciprocity etc.).

One might for example argue, that some participants might have preferences over the outcomes for other participants in the treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$, but not over the computer's outcomes in treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ and that this affects our results. To address this potential concern we ran a treatment that coincides with $\mathbf{E}_{asy}\mathbf{T}_{riv}$, but where computers are replaced by human opponents. We are interested in seeing whether this change reverses any of our treatment rankings and therefore affects any of the conclusions we drew.

	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$-$h$
Entropy	3.21	2.79	2.45	1.13
sd	3.27	2.42	3.15	0.91

Table 5: Behavioral variation across treatments. Data from incentivized rounds in the the second half of the experiment.

Table 5 corresponds to Table 2, but treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$has been replaced by treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$-$h$. It can be seen that the conclusions are the same. Irrespective of the measure (entropy or standard deviation) behavioural variation is much lower in $\mathbf{E}_{asy}\mathbf{T}_{riv}$-$h$ compared to any of the other treatments.

Table 6 corresponds to Table 3, but again treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$ has been replaced by treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$-$h$. Again our conclusions are robust. The share of choices corresponding to best response behaviour or being "close" to best response behaviour is much larger in $\mathbf{E}_{asy}\mathbf{T}_{riv}$-$h$

	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$-**h**
$P(x = NE)$	7.04%	13.33%	-	-
$P(x = BR)$	7.04%	12.04%	22.50%	50.42%
$P(\|x - NE\| \leq 1)$	25.74%	32.78%	-	-
$P(\|x - BR\| \leq 1)$	26.30%	31.30%	47.95%	74.58%
$P(x > 4)$	60.19%	62.78%	51.36%	23.33%

Table 6: Indicators measuring the explanatory power of Nash equilibrium across treatments. Data from the incentivized rounds in the the second half of the experiment.

compared to any of the other treatments. Hence, while some participants may well have social preferences we conclude that (i) homogeneous social preferences cannot explain the large (within subject) variation in behaviour we observe and (ii) our treatment rankings and conclusions are not affected by allowing for this possibility. In the next subsection we ask whether *heterogeneity* in social or other preferences can explain the large behavioural variation we observe.

5.3 Heterogeneity

In this subsection we report the results of simulations based on a number of different populations with heterogeneous preferences. Since in our experiment (treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$) we had matching groups of 6 players, we form arbitrary populations consisting of 6 players with different preferences. The preference types we consider are the following

- τ_1: risk neutral expected utility maximizers.
- τ_2: risk averse CRRA agents with CRRA parameter $r = 0.5$ (see Appendix A)
- τ_3: risk-seeking CRRA agents with CRRA parameter $r = -0.5$ (see Appendix A)
- τ_4: social preferences: Charness-Rabin $(\rho, \sigma) = (0.4, 0)$ (Charness and Rabin (2002))
- τ_5: social preferences: Charness-Rabin $(\rho, \sigma) = (0.8, 0.1)$ (Charness and Rabin (2002))
- τ_6: joy of winning preferences: additional utility of 8 for winning (Sheremeta (2013))

Types $\tau_1 - \tau_3$ vary risk preferences. Types $\tau_4 - \tau_5$ very social preferences as in Charness and Rabin (2002). Type τ_4 reflects the parameters estimated from their experiments and type τ_5 has the same parameters but multiplied by two. Type τ_6 receives an additional utility of 8 if s/he wins the contest. These type of "joy of winning" preferences have received a lot of attention in the contest literature to explain the so-called "overbidding phenomenon" (see Sheremeta (2013)). An additional utility of 8 reflects a 50% increase in expected utility at the risk-neutral Nash equilibrium.

We simulated different population compositions of these types for 40 periods. Before the first period we randomly match agents into pairs. In period 1 they all play an initial action drawn uniformly at random from $1, ..., 16$. In all subsequent periods they play a myopic best response to the action of their previous match. To compute the myopic best response for the Charness-Rabin types τ_4 and τ_5 we assume that the prize is shared proportionally as in treatment $\mathbf{E}_{asy}\mathbf{D}_{iff}$. All populations compositions are run 100 times to account for path dependence caused by random matchings and initial actions. We report the minimal, maximal, mean and median entropy across these 100 runs.

composition $(\tau_1, ..., \tau_6)$	min	max	median	mean
$(6, 0, 0, 0, 0, 0)$	0.217	0.530	0.391	0.379
$(0, 6, 0, 0, 0, 0)$	0.255	0.538	0.430	0.425
$(0, 0, 6, 0, 0, 0)$	0.147	0.429	0.321	0.315
$(0, 0, 0, 6, 0, 0)$	0.806	2.104	1.624	1.656
$(0, 0, 0, 0, 6, 0)$	0.754	2.076	1.613	1.584
$(0, 0, 0, 0, 0, 6)$	0.213	0.426	0.333	0.328
$(0, 2, 0, 2, 2, 0)$	0.320	0.974	0.710	0.694
$(1, 1, 1, 1, 1, 1)$	1.158	1.649	1.415	1.409
$(1, 2, 0, 1, 1, 1)$	1.339	1.695	1.550	1.543
$(1, 2, 0, 2, 1, 0)$	0.327	0.541	0.549	0.758
$(2, 0, 0, 2, 1, 1)$	1.099	1.587	1.330	1.344
$(2, 0, 0, 2, 2, 0)$	0.406	0.952	0.690	0.683
$(2, 1, 1, 1, 1, 0)$	0.252	0.628	0.444	0.447
$(2, 2, 0, 1, 1, 0)$	0.262	0.638	0.464	0.464
$(3, 0, 0, 2, 1, 0)$	0.247	0.790	0.547	0.536
$(3, 0, 0, 3, 0, 0)$	0.192	0.801	0.522	0.522
$(3, 1, 0, 1, 1, 0)$	0.264	0.652	0.473	0.466
$(3, 2, 0, 1, 0, 0)$	0.255	0.597	0.425	0.423
$(3, 3, 0, 0, 0, 0)$	0.186	0.525	0.412	0.405
$\mathbf{D}_{iff}\mathbf{D}_{iff}$				**3.21**
$\mathbf{E}_{asy}\mathbf{D}_{iff}$				**2.79**

Table 7: Minimal, maximal, median and mean entropy across 100 runs of simulated populations. Mean entropy in treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$and $\mathbf{E}_{asy}\mathbf{D}_{iff}$from incentivized rounds in second half of experiment.

Table 7 shows the results of our simulations. The highest levels of entropy are reached in homogeneous populations of agents with Charness-Rabin preferences. The reason is that, because there are multiple Nash equilibria in this case, the process converges much more slowly. However, even in these populations the average levels of entropy observed (1.656 or 1.584, respectively) fall well short of the high levels of entropy observed in treatment $\mathbf{D}_{iff}\mathbf{D}_{iff}$(where mean entropy equals 3.21). We conclude that heterogeneity in preferences cannot explain the high behavioural variation observed in treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$, $\mathbf{E}_{asy}\mathbf{D}_{iff}$and $\mathbf{D}_{iff}\mathbf{T}_{riv}$.

6 Conclusion

We conducted an experiment to understand the reasons behind the large behavioural variation and the low explanatory power of Nash equilibrium typically found in Tullock contests. Across treatments we vary the difficulty of (i) forming correct beliefs and (ii) formulating best responses. In the treatment where both belief formation and best response formulation are "easy", behavioural variation is substantially lower and the explanatory power of Nash equilibrium much higher. Via additional treatments and several simulations we show that heterogeneity in preferences cannot explain the large behavioural variation found in the standard contest. We conclude that bounded rationality rather than preference heterogeneity is the reason for the typically large behavioural variation in experimental Tullock contests.

References

Abbink, K., J. Brandts, B. Herrmann, and H. Orzen (2010). Intergroup conflict and intra-group punishment in an experimental contest game. *The American Economic Review 100*(1), 420–447.

Andersson, M. and Y. Iwasa (1996). Sexual selection. *Trends in Ecology and Evolution 11(2)*, 53–58.

Baharad, E. and S. Nitzan (2008). Contest efforts in light of behavioural considerations. *The Economic Journal 118*(533), 2047–2059.

Bartling, B., E. Fehr, M. Marechal, and D. Schunk (2009). Egalitarianism and competitiveness. *American Economic Review 99*, 93–98.

Baye, M., D. Kovenock, and C. D. Vries (1993). Rigging the lobbying process: An application of the all-pay auction. *The American Economic Review 83(1)*, 289–294.

Baye, M. R. and H. Hoppe (2003). The strategic equivalence of rent-seeking, innovation and patent-race games. *Games and Economic Behavior 44(2)*, 217–226.

Bednar, J., Y. Chen, T. Liu, and S. Page (2011). Behavioral spillovers and cognitive load in multiple games: An experimental study. *Games and Economic Behavior*.

Brosig, J. and J. Reiss (2007). Entry decisions and bidding behavior in sequential first price procurement auctions: An experimental study. *Games and Economic Behavior 58(1).*

Buchanan, J. and G. Tullock (1962). *The Calculus of Consent: Logical Foundations of a Constitutional Democracy.* Ann Arbor: University of Michigan Press.

Cason, T., W. Masters, and R.M.Sheremeta (2010). Entry into winner-take-all and proportional-prize contests. *Journal of Public Economics 94*, 604–611.

Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. *American Economic Review 117(3)*, 817–869.

Chen, K. (2003). Sabotage in promotion tournaments. *Journal of Law, Economics and Organization 19(1)*, 119–140.

Cooper, D. and H. Fang (2008). Understanding overbidding in second price auctions: An experimental study. *Economic Journal 118*, 1572–1595.

Esteban, J. and D. Ray (2011). A model of ethnic conflict. *The Journal of the European Economic Association 9*, 496–521.

Fallucchi, F., E. Renner, and M. Sefton (2013). Information feedback and contest structure in rent seeking games. *European Economic Review forthcoming.*

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

Greiner, B. (2004). The online recruitment system orsee 2.0-a guide for the organization of experiments in economics. *University of Cologne, Working paper series in economics 10*(23), 63–104.

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

Hillman, A. L. and E. Katz (1984). Risk-averse rent seekers and the social cost of monopoly power. *The Economic Journal 94*(373), 104–110.

Huck, S., P. Jehiel, and T. Rutter (2010). Feedback spillover and analogy based expectations: A multi game experiment. *Games and Economic Behavior 71(2)*, 351365.

Huck, S. and G. Weizsäcker (1999). Risk, complexity, and deviations from expected-value maximization: Results of a lottery choice experiment. *Journal of Economic Psychology 20*(6), 699–715.

Johnson, E., C. Camerer, S. Sen, and T. Rymon (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory 104*(1), 16–47.

Kahneman, D. and A. Tversky (1972). Subjective probability: A judgement of representativeness. *Cognitive Psychology 3(3)*, 430–454.

Konrad, K. (2009). *Strategy and Dynamics in Contests.* Oxford University Press.

Lim, W., A. Matros, and T. Turocy (2014). Bounded rationality and group size in tullock contests. *Journal of Economic Behavior and Organization forthcoming.*

Millner, E. and M. Pratt (1989). An experimental investigation of efficient rent-seeking. *Public Choice 62*(2), 139–151.

Peters, E., N. Dieckmann, A. Dixon, J. Hibbard, and C. Mertz (2007). Less is more in presenting quality information to consumers. *Medical Care Research and Review 64*(2), 169–190.

Potters, J., C. De Vries, and F. Van Winden (1998). An experimental examination of rational rent-seeking. *European Journal of Political Economy 14*(4), 783–800.

Rabin, M. and G. Weizsäcker (2009). Narrow bracketing and dominated choices. *American Economic Review 99*, 1508–1543.

Schmitt, P., R. Shupp, K. Swope, and J. Cadigan (2004). Multi period rent-seeking contests with carryover: Theory and experimental evidence. *Economics of Governance 5*, 187–211.

Shannon, C. (1948). The bell system technical journal 27. *A Mathematical Theory of Communication*, 379–423.

Sheremeta, R. (2010). Experimental comparison of multi-stage and one-stage contests. *Games and Economic Behavior 68*(2), 731–747.

Sheremeta, R. (2011). Contest design: an experimental investigation. *Economic Inquiry 49*(2), 573–590.

Sheremeta, R. (2013). Overbidding and heterogeneous behavior in contests. *Journal of Economic Surveys 27*, 491–514.

Szymanski, S. (2003). The economic design of sporting contests. *Journal of Economic Literature 41*, 1137–1187.

A Additional Tables and Figures

A.1 Risk Preferences

Figure 6 illustrates the best response correspondence for an agent with CRRA preferences $U(x) = \frac{x^{1-r}}{1-r}$ with parameter $r \in \{-0.5, 0, 0.5\}$. The figure shows that all choices above 4 are dominated for all these types. Best response correspondences are of similar shape and the Nash equilibrium is $(4, 4)$ for any pairwise match between agents of either of these three types.

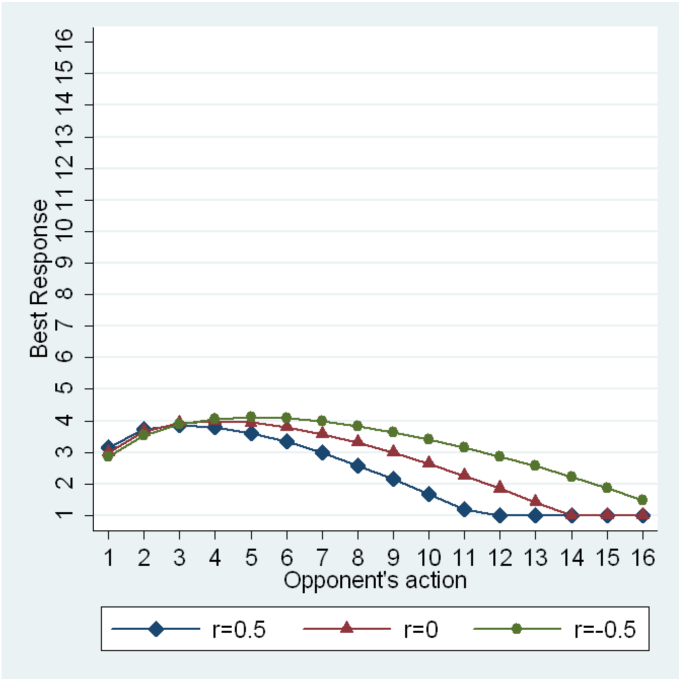

Figure 6: Best Response Correspondence for an agent with CRRA preferences with parameter $r \in \{-0.5, 0, 0.5\}$

A.2 Regression Explanatory Power of Nash equilibrium

Table 8 shows the results of logit regressions where either (in columns (1) and (2)) we regress a binary variable indicating an undominated choice ($a < 5$) on treatment dummies or (in columns (3) and (4)) we regress a binary variable indicating a best response to the opponent's choice on treatment variables. The results illustrate the statistical significance of the differences seen in Table 3. A χ^2 test shows that the difference between $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}$ is smaller than the difference between $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ ($p < 0.0023$).

	$\Pr(a<5)$	$\Pr(a<5)$	$\Pr(a \in BR(a_{-}i))$	$\Pr(a \in BR(a_{-}i))$
"E_{asy}..."	-0.109	1.702^{***}		
	(0.392)	(0.328)		
"...T_{riv}"			1.344^{***}	2.617^{***}
			(0.234)	(0.416)
constant	-0.413^{**}	-0.054	-2.581^{***}	-1.988^{***}
	(0.216)	(0.198)	(0.130)	(0.355)
Observations	1080	880	980	980
Baseline	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Logit Regressions Explanatory Power of Nash equilibrium. Standard errors clustered at matching group level. The first column compares the frequency of undominated choices in treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$, the second column compares $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$. Column (3) compares the share of choices that are best responses to the opponent's behaviour in $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and column (4) compares $\mathbf{E}_{asy}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$.

A.3 Deviations from best responses across blocks

Table 9 shows how much participants choices differ from best responses in each block of the game, how big the effect of removing *either* source of uncertainty is and whether these effects are statistically significant in each block. The statistical comparison between $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$ treatments was done using matching groups as independent observations, while the comparison between $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$ was done using individual players as independent observations. If individual players are used for the first comparison instead, results do not change and p-values are above 0.1 in every block. Student's t test yields the similar results, except that the difference between means of treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$ is marginally significant in blocks 2, 3 and 4, respective p-values 0.0633, 0.0415 and 0.0186.

Treatment	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$
First block	3.77	3.04	4.45	1.59
Second block	3.03	2.60	4.41	1.15
Third block	2.86	2.53	3.69	0.74
Fourth block	2.90	2.43	3.23	0.60
Total	3.14	2.65	3.95	1.02

Comparison	$\mathbf{D}_{iff}\mathbf{D}_{iff}$ vs $\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$ vs $\mathbf{E}_{asy}\mathbf{T}_{riv}$
First block	0.0094	0.0000
Second block	0.3008	0.0000
Third block	0.1050	0.0000
Fourth block	0.1224	0.0000
Total	0.0014	0.0000

Table 9: Comparison of mean equilibrium deviations across treatments in every block. Two-tailed Mann-Whitney U test p-values of pairwise comparisons in the right panel

Table 10 shows the median investment for each strategy played by the opponent in treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$. The table illustrates that the strong deviations from best response behaviour in $\mathbf{D}_{iff}\mathbf{T}_{riv}$ are not due to outliers but are a systematic pattern. The table also shows that median choices are largely consistent with best response behaviour in $\mathbf{E}_{asy}\mathbf{T}_{riv}$.

Strategy	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Theoretical BR	3	4	4	4	4	4	4	3	3	3	2	2	1	1	1	1
$\mathbf{E}_{asy}\mathbf{T}_{riv}$(last 20)	3	4	4	4	4	4	4	3	3	3.5	2	2	1	1	1	1
$\mathbf{D}_{iff}\mathbf{T}_{riv}$(last 20)	3	5	4	5	5	8	10	8	9	11	1	1	0.5	7	1	1

Table 10: Theoretical best-responses and the median actual investments for each strategy played by the computer.

A.4 Determinants Behaviour

Table 11 shows the results of a regression looking at what factors induce a player to deviate from the theoretical prediction. Among the explanatory variables we include personal characteristics:

the total number of correct answers in a numeracy test and gender as well as variables on the history of play: opponent's investment in the previous period (treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$), strategy chosen by the computer (treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$), a binary variable indicating whether the agent won in the previous period and the inverse of the period variable.

Results of the regression are presented in Table 11. Players who did better in the numeracy test on average behave more in line with the theoretical prediction in treatment $\mathbf{D}_{iff}\mathbf{T}_{riv}$ (see Section 4.3). No significant gender effect is observed. In all the treatments behavior tends to move towards the equilibrium over time, as seen from the positive coefficient on the inverse period variable, but the learning effect is strongest in treatment $\mathbf{E}_{asy}\mathbf{T}_{riv}$.

	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$
numeracy score	-0.142	-0.133	-0.648***	-0.214*
Female	-0.128	-0.582	0.0426	0.112
1/period	6.370*	5.777**	6.800	7.808***
Won in the previous round	-0.103		0.973***	
Opponent's action in the previous round	0.0598*	0.0584**		
Computer's investment			0.331***	-0.0136
Constant	3.217**	3.021***	3.815***	1.754***
Observations	980	1000	820	820

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 11: Random effects estimation, the dependent variable is absolute difference between the best response prediction and investment level. Standard errors clustered on subject level. Data from all incentivized rounds.

A.5 Convergence and Dynamics

Figure 7 shows the frequency with which participants switch actions over time in the non-incentivized and incentivized periods. The figure provides some interesting insights. In the treatments $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{E}_{asy}\mathbf{D}_{iff}$there seems to be little learning over time. Participants switch their choices with a probability of around 0.6 across periods. This number barely declines over time (with the exception maybe of the first 5 periods, where it drops from 0.8 to 0.6) and there is no discernible difference between incentivized and non-incentivized periods. There seems to be almost no learning in these treatments and participants remain unsure about what to choose until the end of the 40 periods.

The picture looks much different in the treatments $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$. In the incentivized rounds, participants switch their choice with a probability of roughly 0.4 ($\mathbf{D}_{iff}\mathbf{T}_{riv}$) or 0.3 ($\mathbf{E}_{asy}\mathbf{T}_{riv}$) on average. Switching probabilities decrease over time, both within and across rounds. And there is a clear difference between incentivized and non-incentivized rounds. Clearly, participants are using the non-incentivized rounds to experiment and learn about best responses and then apply these in the incentivized rounds. The difference between treatments with difficult and trivial belief formation is highly statistically significant (one-sided t-test $p < 0.0001$), so is the difference between $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{E}_{asy}\mathbf{T}_{riv}$(one-sided t-test on incentivized rounds only

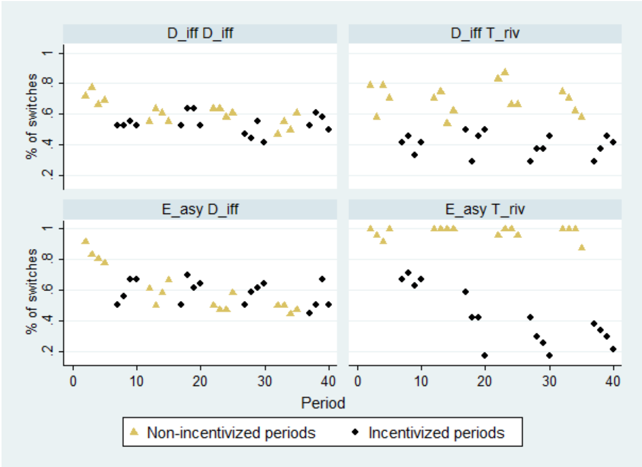

Figure 7: Relative frequency of changes in investments from one round to the other, by treatment.

$p = 0.0392$.

A.6 Within-subject and between-subject variability

Table 12 decomposes the behavioural variation (standard deviation) into within and between subject variation. The table reports both sum of squares and mean of squares each decomposed into between and within subject variation. The table shows that both within-subject and between-subject variability is lower in $\mathbf{E}_{asy}\mathbf{D}_{iff}$ compared to $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and lower in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ compared to $\mathbf{D}_{iff}\mathbf{T}_{riv}$. Hence making it easier to formulate best responses lowers both within and between subject variability irrespective of whether belief formation is difficult or trivial. Furthermore, within subject variation is lower in $\mathbf{E}_{asy}\mathbf{T}_{riv}$ compared to $\mathbf{E}_{asy}\mathbf{D}_{iff}$, but the difference between $\mathbf{D}_{iff}\mathbf{T}_{riv}$ and $\mathbf{D}_{iff}\mathbf{D}_{iff}$is not significant. Overall, $\mathbf{E}_{asy}\mathbf{T}_{riv}$ has the lowest level of both types of variability whereas $\mathbf{D}_{iff}\mathbf{D}_{iff}$ and $\mathbf{D}_{iff}\mathbf{T}_{riv}$ are most variable.

Treatment	**MSQ Within**	**MSQ Between**	**SSQ Within**	**SSQ Between**
$\mathbf{D}_{iff}\mathbf{D}_{iff}$	3.96	17.34	1.65	9.17
$\mathbf{E}_{asy}\mathbf{D}_{iff}$	2.27	11.16	1.14	5.55
$\mathbf{D}_{iff}\mathbf{T}_{riv}$	2.74	49.15	2.28	8.28
$\mathbf{E}_{asy}\mathbf{T}_{riv}$	0.36	7.71	0.30	1.25

Table 12: Within-subject and between-subject variability by treatment. Variables: mean squares (MSQ) and sum of squares (SSQ) within subjects and between subjects. Variability computed conditional on opponent's investment and then averaged using the frequency of each strategy chosen in the standard contest as weights. Data from incentivized rounds in the second half of the experiment.

B Instructions $\mathbf{D}_{iff}\mathbf{D}_{iff}$

INSTRUCTIONS

Welcome to the experiment. Please read the instructions carefully. They are identical for all the participants with whom you will interact during this experiment.

If you have any questions please raise your hand. One of the experimenters will come to you and answer your questions. From now on communication with other participants is not allowed. If you do not conform to these rules you will be excluded from the experiment with no payment. Please do also switch off your mobile phone at this moment.

In this experiment you can earn some money. How much you earn depends on your decisions and the decisions of the other participants. During the experiment we will refer to ECU (Experimental Currency Unit) instead of Euro. The total amount of ECU that you will have earned during the experiment will be converted into Euro at the end of the experiment and paid to you in cash confidentially. The conversion rate that will be used to convert your ECU earnings into your Euro cash payment will be shown to you on the screen at the beginning of the experiment.

The Experiment

The main part of this experiment consists of 4 blocks with 10 rounds in each block. In each block, the first five rounds are for practice only so that you can experiment without affecting your cash earnings from this experiment. At the end of the experiment, one round out of the final five rounds in each block (that is one round of rounds 6-10 in each block) will be randomly selected and the sum of your round incomes in these selected rounds will be converted into euros and paid to you in cash.

The Task

The task is the same in each of the 40 rounds. At the beginning of each round the computer will randomly match you with another participant in this room. You will not know who the other participant is, and the other participant you are matched with is likely to change every round. The other participant that is matched with you will receive the same information and will face exactly the same task. In each round each participant will receive an endowment of 16 ECU. The endowment can be used to purchase "tokens". Each token costs 1 ECU so that you can purchase up to 16 of these tokens. You have to buy at least one token. Any part of your endowment that you do not spend on tokens will be added to your round income.

After you and the other participant have chosen how many tokens to buy, only one of you will receive the extra 16 ECU. The probability that you will receive the extra 16 ECU depends on the number of tokens that you buy and the number of tokens that the other participant buys. More precisely, the probability that you receive the extra 16 ECU is given by:

$$\begin{array}{l}\text{Probability of receiving}\\ \text{the extra 16 ECU}\end{array} = \frac{\text{Number of tokens you bought}}{\begin{array}{l}\text{Number of tokens you bought + Number}\\ \text{of tokens the other participant bought}\end{array}} \times 100\%$$

For example, if both of you have purchased the same number of tokens, the probability that each of you will receive the extra 16 ECU is 50%. Note that either you or the other participant will always receive the extra 16 ECU.

Whether you or the other participant will receive the extra 16 ECU will be determined by a random draw by the computer according to the probabilities given by the number of tokens bought by you and by the other participant. Then the computer will compute your round income based on the number of tokens that you bought and whether you have received the extra 16 ECU or not. Once the round is over, you will be informed about the number of tokens bought by you and by the other participant, the probability to receive the extra 16 ECU, the outcome of who receives the extra 16 ECU and your round income, and that of the other participant.

- If you receive the extra 16 ECU, your round income will be:

  $$\text{Round Income} = 16\text{ ECU} - \text{Number of purchased tokens} + 16\text{ ECU}$$

- If you do not receive the extra 16 ECU, your total earnings in the round will be:

  $$\text{Round Income} = 16\text{ ECU} - \text{Number of purchased tokens}$$

At the end of the experiment, four rounds will be randomly selected for payment. More precisely, the first round for payment will be randomly selected from rounds 6-10, the second round for payment from rounds 16-20, the third from rounds 26-30 and the fourth one from rounds 36-40. Outcomes in all other rounds will not influence your final earnings, but you will not know which rounds will be selected until the end of the experiment.

After the 40 rounds we will ask you to fill in a short questionnaire. The questionnaire will have two parts, each of which will be explained on the screen before you start answering the questions. After the first part of the questionnaire you will be informed about all of your round incomes as well as about the four rounds that were randomly selected for payment (see Figure 1 on the next page). After the second part of the questionnaire you will be informed about your final earnings in euro. You will receive these earnings in cash and in private at the end of the experiment. Please stay seated until we ask you to come to receive the earnings.

If you have any further questions, please raise your hand now.

Summary

The structure of the experiment is as follows:

- The main part of the experiment consisting of 40 rounds.
- Questionnaire, part 1. After answering these questions you will be informed about your final earnings in ECU.
- Questionnaire, part 2. Once you have completed it, you will be informed about your final cash earnings in euros.
- Please stay seated until the experimenter asks you to come and receive the earnings.

C Instructions $\mathbf{E}_{asy}\mathbf{T}_{riv}$

INSTRUCTIONS

Welcome to the experiment. Please read the instructions carefully. They are identical for all the participants with whom you will interact during this experiment.

If you have any questions please raise your hand. One of the experimenters will come to you and answer your questions. From now on communication with other participants is not allowed. If you do not conform to these rules you will be excluded from the experiment with no payment. Please do also switch off your mobile phone at this moment.

In this experiment you can earn some money. How much you earn depends on your decisions and the decisions of the other participants. During the experiment we will refer to ECU (Experimental Currency Unit) instead of Euro. The total amount of ECU that you will have earned during the experiment will be converted into Euro at the end of the experiment and paid to you in cash confidentially. The conversion rate that will be used to convert your ECU earnings into your Euro cash payment will be shown to you on the screen at the beginning of the experiment.

The Experiment

The main part of this experiment consists of 4 blocks with 10 rounds in each block. In each block, the first five rounds are for practice only so that you can experiment without affecting your cash earnings from this experiment. At the end of the experiment, one round out of the final five rounds in each block (that is one round of rounds 6-10 in each block) will be randomly selected and the sum of your round incomes in these selected rounds will be converted into euros and paid to you in cash.

The Task

The task is the same in each of the 40 rounds. In each round you will receive an endowment of 16 ECU. The endowment can be used to purchase "tokens". Each token costs 1 ECU so that you can purchase up to 16 of these tokens. You have to buy at least one token. Any part of your endowment that you do not spend on tokens will be added to your round income.

In every round you will be matched to a computerized participant (computer). The computer will buy a certain number of tokens (between 1 and 16), and this number is pre-determined before the start of the experiment. The number of tokens bought by the computer in any round will be announced on the screen before you make your buying decision for that round. This amount will be the same in each round of a block, but will change from one block to another. That means that the computer will buy the same number of tokens in rounds 1-10, 11-20, 21-30 and 31-40.

After you have chosen how many tokens to buy, either you or the computer will receive an extra 16 ECU. The probability that you will receive the extra 16 ECU depends on the number of tokens that you buy and the number of tokens that are bought by the computer. More precisely, the probability that you receive the extra 16 ECU is given by:

$$\text{Probability of receiving the extra 16 ECU} = \frac{\text{Number of tokens you bought}}{\text{Number of tokens you bought + Number of tokens the computer bought}} \times 100\%$$

For example, if you purchase the same number of tokens as the computer, the share of 16 ECU will be equal to 50%, meaning that you will receive 8 ECU.

Your round income will be computed based on the number of tokens that you bought and the share of the extra 16 ECU that you have received. Once the round is over, you will be informed about the number of tokens bought by you and by the computer, the share of the extra 16 ECU, your round income, and that of the computer. Information about the number of tokens and the share of 16 ECU allocated to you and to the other participant will also be represented visually.

- Your round income will be:

  Round Income = 16 ECU − Number of purchased tokens + (Share of the extra 16 ECU)*16 ECU

At the end of the experiment, four rounds will be randomly selected for payment. More precisely, the first round for payment will be randomly selected from rounds 6-10, the second round for payment from rounds 16-20, the third from rounds 26-30 and the fourth one from rounds 36-40. Outcomes in all other rounds will not influence your final earnings, but you will not know which rounds will be selected until the end of the experiment.

After the 40 rounds we will ask you to fill in a short questionnaire. The questionnaire will have two parts, each of which will be explained on the screen before you start answering the questions. After the first part of the questionnaire you will be informed about all of your round incomes as well as about the four rounds that were randomly selected for payment (see Figure 1 on the next page). After the second part of the questionnaire you will be informed about your final earnings in euro. You will receive these earnings in cash and in private at the end of the experiment. Please stay seated until we ask you to come to receive the earnings.

If you have any further questions, please raise your hand now.

Summary

- The main part of the experiment consisting of 40 rounds.
- Questionnaire, part 1. After answering these questions you will be informed about your final earnings in ECU.
- Questionnaire, part 2. Once you have completed it, you will be informed about your final cash earnings in euros.
- Please stay seated until the experimenter asks you to come and receive the earnings.

Your income in the payoff-relevant rounds

Block 1		Block 2		Block 3		Block 4	
Round	Income	Round	Income	Round	Income	Round	Income
6	22.00	16	25.00	26	28.00	36	24.00
7	31.00	17	4.00	27	25.00	37	17.00
8	26.00	18	31.00	28	25.00	38	4.00
9	10.00	19	20.00	29	7.00	39	4.00
10	11.00	20	26.00	30	26.00	40	26.00

The rounds that were randomly selected for payment are:

10, 19, 26 and 39.

Figure 8: A screenshot of the final profit display. Round incomes in this screenshot were randomly generated. [Picture used in the instructions of all treatments]

D Questionnaire

D.1 Questions

Apart from questions about age, gender, bachelor programme and previous experience in the lab, the following 8 questions were asked in the numeracy questionnaire.

- If Person As chance of getting a disease is 1 in 100 in 10 years, and person Bs risk is double that of A, what is Bs risk?
- Imagine that you are taking a class and your chances of being asked a question in class are 1% during the first week of class and double each week hereafter (i.e., you would have a 2% chance in Week 2, a 4% chance in Week 3, an 8% chance in Week 4). What is the probability that you will be asked a question in class during Week 7? (in %)
- Suppose that 1 out of every 10,000 doctors in a certain region is infected with the SARS virus; in the same region, 20 out of every 100 people in a particular at-risk population also are infected with the virus. A test for the virus gives a positive result in 99% of those who are infected and in 1% of those who are not infected. A randomly selected doctor and a randomly selected person in the at-risk population in this region both test positive for the disease. Who is more likely to actually have the disease?
- In the Acme Publishing Sweepstakes, the chance of winning a car is 1 in 1,000. What percentage of tickets of Acme Publishing Sweepstakes wins a car?
- Imagine that we roll a fair, six-sided die 1,000 times. Out of 1000 rolls, how many times do you think the die would come up even (2, 4, or 6)?

- In the Big Bucks Lottery, the chances of winning a 10 prize are 1%. What is your best guess about how many people would win a 10 prize if 1000 people each buy a single ticket from Big Bucks?
- The chance of getting a viral infection is 0.0005. Out of 10,000 people, about how many of them are expected to get infected?
- Suppose you have a close friend who has a lump in her breast and must have a mammogram. Of 100 women like her, 10 of them actually have a malignant tumor and 90 of them do not. Of the 10 women who actually have a tumor, the mammogram indicates correctly that 9 of them have a tumor and indicates incorrectly that 1 of them does not. Of the 90 women who do not have a tumor, the mammogram indicates correctly that 81 of them do not have a tumor and indicates incorrectly that 9 of them do have a tumor. The graph below summarizes all of this information. Imagine that your friend tests positive (as if she had a tumor), what is the likelihood that she actually has a tumor?

D.2 Summary Statistics

	$\mathbf{D}_{iff}\mathbf{D}_{iff}$	$\mathbf{E}_{asy}\mathbf{D}_{iff}$	$\mathbf{D}_{iff}\mathbf{T}_{riv}$	$\mathbf{E}_{asy}\mathbf{T}_{riv}$
age	21.86	21.57	21.70	21.85
	(18,41)	(18,26)	(18,26)	(19,30)
female	0.36	0.53	0.45	0.63
	(0,1)	(0,1)	(0,1)	(0,1)
numeracy score	5.20	5.22	5.11	5.27
	(1,7)	(0,8)	(1,8)	(0,8)

Table 13: Summary statistics from questionnaire (Mean and Range). The variable female takes the value 1 for women. The variable numeracy score counts the number of correct answers in the numeracy test.