(De)-Anchoring Beliefs in Beauty Contest Games*

Jess Benhabib
NYU

John Duffy†
UCI

Rosemarie Nagel
ICREA
UPF
Barcelona GSE

November 6, 2019

Abstract

The beauty contest game (BCG) serves as a core framework for demonstrating behavioral non-equilibrium patterns such as focal points and level-k reasoning. We introduce a new version of the BCG that removes the bounded choice interval and thus eliminates iterative elimination of dominated strategies. We further add correlated idiosyncratic signals that can serve as (equilibrium) coordination devices. We find that choices in these new versions of the BCG are closer to equilibrium as compared with the standard BCG. Indeed, we show how variations in the design of BCGs can greatly affect the use of focal points and level-k reasoning.

JEL codes: C92, D83, D84.
Keywords: Beauty contest game, expectation formation, equilibration, level-k reasoning, macroeconomics, game theory, experimental economics.

*We thank Larbi Alaoui, Antonio Cabrales, David Levine, Sheen Levine, Antonio Penta, and Michael Woodford for many helpful comments. Special thanks go to Roger Guesnerie for organizing a workshop on coordination the INECX program in Paris 2012, where the idea for the paper was born. We are also grateful for comments received from participants at several other conferences and workshops. Thomas Stoerk and Felix Mauersberger provided expert research assistance under FPI grant no. ECO2014-56154-P. Nagel gratefully acknowledges the hospitality of NYU-CES Department of Economics during a sabbatical year. She further acknowledges funding from a Spanish Ministry of Education grant mobility grant; Spanish Ministry of Education grants ECO2011-25295, ECO2014-56154-P, Generalitat de Catalunya (Agaur), and Barcelona Graduate School of Economics. Duffy gratefully acknowledges funding for this project from the Dietrich School of Social Sciences at the University of Pittsburgh.

†Corresponding author: duffy@uci.edu

"human decisions affecting the future, whether personal, political or economic, cannot depend on strict mathematical expectation, since the basis for making such calculations does not exist.... [In making decisions] our rational selves choose between the alternatives as best as we are able, calculating where we can, but often falling back for our motive on whim or sentiment or chance." [Keynes (1936), pp. 162-163]

1 Introduction

People often make suboptimal choices due to anchoring, salience or limited cognitive or strategic sophistication. The Keynesian beauty contest game (Keynes 1936, Moulin 1983), is a core framework for examining and revealing such behavior, using the so-called level-k reasoning model, which reveals the degree of a player's iterated best reply to an initial intuitive or random prior belief about others (level 0). In the basic game, n players privately and simultaneously choose a number from some closed interval, typically 0 to 100 inclusive, with the objective that their number is closest in absolute value to β times the average of all n chosen numbers, with e.g., $\beta = 2/3$.

The Keynesian Beauty contest game *also* serves as a kind of reduced form best reply function for a large class of games and models both in the micro- and macro-economic literature, where higher order beliefs play a role. Models in this class differ primarily in the choice of the "tuning" parameter, β, which can be positive (strategic substitutes) or negative (strategic complements), but also differ in whether the strategy space is discrete or continuous, or in the informational structure, e.g., whether agents have access to noisy signals or not.[1]

Our aim in this paper is to explore belief formation in a *new* version of the beauty contest (BC) game that is a simplified version of a general equilibrium model with or without sentiments (signals) that is due to Benhabib et al. (2015). This new version can be thought of a BC game with an unbounded guessing interval, where correlated random signals (or sentiments) can rationally serve as a coordinating device. This version of the BC game enables us to better understand the (de)-anchoring and adjustment of beliefs to information than does the standard version of the BC game. While this new version of the BC game has little to no impact on theoretical predictions relative to the standard BC game, the changes made to the game nevertheless have a large impact on actual behavior, and lead us to the more optimistic conclusion that, depending on the setting, agents may be able to very quickly solve the coordination problem associated with forming beliefs about the beliefs of others.

In our new version of the BC game, we first eliminate the bounded guessing interval and instead allow subjects to choose *any* number on the real line (negative, positive or zero). In many economic coordination problems, the strategy space is similarly unbounded, e.g., profits or rates of inflation can be positive, negative or zero. We show that the elimination of boundaries for the guessing interval effectively "de-anchors" subjects' initial beliefs relative to BC games with a bounded guessing interval and $\beta \in (0, 1)$. A result of this de-anchoring is that subjects' guesses are

[1] Camerer and Fehr (2006) classify some experimental games according to strategic substitutes vs. complements, and discrete vs. continuous strategy spaces, with the original BC game being just one of the games with strategic complements; Mauersberger and Nagel (2018) show how a general, beauty contest game best reply function as introduced in the appendix characterizes many archetypal micro and macro models, e.g., the Cournot model, ultimatum bargaining games, public good games, some auctions, New Keynesian models, asset pricing models, etc.). Angeletos and Lian (2016) discuss the theoretical effects of different information structures distinguishing between global games (normal form games) and BC-games.

initially closer to the unique equilibrium of the BC game than in the standard, bounded guessing interval case. Second, in some treatments, we introduce noisy signals (or sentiments) that are either perfectly correlated with one another or have a common component. We believe that such signals better approximate the environments in which individuals are called upon to form their beliefs. In one of these signal settings, the idiosyncratic information provided to each individual can be interpreted as a correlated signal of the true state of the world (or variable) and predicting this state is the common goal of individual expectation formation. By contrast, in the standard BC game, no such signals are provided. In the case where signals have a common and idiosyncratic component, subjects have to engage in a signal extraction exercise as in the Lucas island model (Lucas 1972).

Third, we replace the traditional tournament payoff structure with a quadratic distance payoff function so that every subject can get a non-zero payoff.[2] With a distance payoff function, individual efficiency calculations become possible which is *not* the case with tournament payoffs. Earlier experiments have shown that this change has no behavioral consequences (see e.g., Gueth et al. 2002 and Nagel et al. 2017).) Subjects' payoffs in our first baseline, no signal treatment depend on the distance between their choice and the forecast target, β times the average of all guesses submitted. In our second and third treatments involving correlated signals and correlated signals with a common component, we add to this target objective an idiosyncratic payoff relevant signal.

As in a correlated equilibrium, it is important that subjects know the distribution of correlated signals. In this particular case, signals are drawn from an i.i.d. mean zero distribution and this fact is known to all. Strictly speaking, however, in a correlated equilibrium, signals should not be payoff relevant, as they are in our signal treatments.[3] Our setup is a simplified version of Benhabib et al. (2015), who study a general equilibrium macroeconomic model with idiosyncratic as well correlated sentiment shocks transmitted to agents via signals who then must solve signal extraction problems to choose their optimal actions. They show that two rational expectation equilibria exist: a certainty equilibrium driven by fundamentals and a stochastic equilibrium (or sentiment-driven equilibrium) in which average production is based on non-fundamental sentiments or self-fulfilling beliefs that are endogenously determined in equilibrium; see the appendix for the details.[4]

It is of interest to see how people behave in these noisy signal environments, as they might well serve as (equilibrium) coordination devices or anchors. Thaler and Sunstein's (2008) examples of nudges are typically for individual decision-making, similar to Tversky and Kahneman's (1986) discussion of anchoring or reference points. The literature speaks less of anchoring in larger, n-player non-cooperative games or even about de-anchoring. To be precise, we regard anchoring as the tendency to rely too heavily on some arbitrary but salient piece of information –"the anchor" – when making decisions. For example, the announced guessing interval of $[0, 100]$ in the standard BC game typically induces an anchor of 50, the midpoint of that guessing interval, as the initial, "level-0"

[2]This design change converts the unique equilibrium of the standard BC game into a unique Pareto efficient equilibrium arrived at by iterated elimination of *strictly* dominated strategies when there are boundaries. Both features are better from both an economic and epistemic point of view.

[3]For this reason, we may wish to refer to the equilibria in our games with signals as Bayes-Nash equilibria rather than correlated equilibria. However, see Bergemann and Morris (2016), who connect standard asymmetric information games having payoff relevant signals with correlated equilibrium and term the resulting construct "Bayes-Correlated" equilibria.

[4]In the full-fledged general equilibrium model, consumers maximize their utility based on wage and price expectations, while firms maximize profits. Consumers send demand signals to the firms, who use these signals for their production decisions, since they have to produce before prices are realized in the goods market.

guess. "De-anchoring" refers to the removal of such anchors. Yet, our new precise signal treatment correctly anchors behavior to the signals as it should in equilibrium. In the more complicated signal extraction treatment, the most prominent heuristic used by subjects is to ignore the signal altogether and play zero, as in the equilibrium of Benhabib et al. (2015) where there are no sentiments.

The past 20 years of experimental research on this version of the BC game (see literature reviews Nagel 1998, Camerer 2003, Crawford et al. 2013, Nagel et al. (2017), Mauersberger and Nagel (2018)), have shown that subjects anchor their initial beliefs about the choices of others far from the equilibrium, and iteratively best respond to those anchors, as explained by the level-k model, (Nagel, 1995, Stahl and Wilson (1995)), and the related cognitive hierarchy model Camerer, Ho, and Chong (2004).[5] The depth of reasoning that individuals exhibit in beauty contest games need not depend on their cognitive abilities alone, but may also be an endogenous function of the strategic environment they face including the payoffs and the sophistication of their opponents, a notion that has been formalized by Alaoui and Penta (2016, 2017) using a cost-benefit approach. Behavior differs greatly within a population, say students, but also between different subject pools such as economics professors or the general public reached through newspaper contests (Bosch et al. 2002).

Fehr et al. (2018) is the closest study related to Keynes (1936) and to our work, yet it differs in important aspects. They study the effect of private and public (noisy), non-payoff relevant signals in a [0,100] choice set, in which 2 players need to coordinate on the same number, thus $\beta = 1$. When there are no signals, subjects choose the midpoint, a focal point, one of multiple equilibria. However, when (noisy) non-payoff relevant signals are added, they work best, when they are public and salient or when private signals are highly correlated. A mixture of private and public signals typically leads to miscoordination. The main focus of that paper is how different information structures of signals induce coordination or miscoordination.[6]

Morris and Shin (2002), formulate a β-beauty contest game with an additional constant term to the target as in our design. However, in their approach, this additional constant term represents a common fundamental value, unknown to the players. Additionally, players receive private and public signals about this fundamental value as in a common value auction. Experimental results on Morris Shin (2002) show that subjects essentially chose some convex combinations of the private and public signal, which typically is far from the inefficient unique equilibrium (see e.g. Baeriswyl and Cornand,(2014)).

Convergence to equilibrium in BC games is achieved mainly by repeating the game several times with the same players, with the speed of convergence depending on the parameter β, the order statistic (e.g. mean vs. median), or the number of players in a group, teams etc. (see Camerer (2003), Nagel et al. (2017). Sutan and Willinger (2009) add a constant term to the guessing target (non-zero equilibria) and study positive and negative feedback versions ($\beta > 0$ or $\beta < 0$) with a bounded positive choice set. They show that negative feedback systems converge faster. Hommes et al. (2013) do something similar, but add random shocks and do not inform subjects of the affine target function generating realizations (they provide only qualitative information about it) as they are interesting in adaptive learning behavior. They find large dampened or explosive

[5] Garcia-Schmidt and Woodford (2016) and Farhi and Werning (2016) introduce level-k in NK-models and study convergence features.

[6] Arifovic and Duffy (2018) discuss many other games with different signal structures and resulting experimental data.

oscillatory fluctuations around the equilibrium in the positive feedback case, but rapid convergence to equilibrium in the negative feedback case. Our new simple BC games without boundaries and without signals can easily explain these feedback differences. Slonim (2005) shows that disequilibrium behavior persists if new subjects are repeatedly injected into the population in an overlapping generation manner.

By contrast, in our new design which involves an unbounded guessing interval, the mean and variance of the initial guesses are much smaller than in any of the previous experiments, and behavior is similar across different subject pools. Our new BC game with signals can produce mathematically more complicated situations: there is variance in equilibrium behavior (because of noisy signals) and in our imprecise signal treatment there is a complicated signal extraction problem that subjects must solve. However, we will show that efficiency is higher in the two new environments with signals than in the BC game without signals. The signals serve as (equilibrium) coordination devices or in the more complex setting, players may simply ignore signals all together leading to near-zero average choices.

Finally, since the success of new rules or institutions often depends on the previous experience of individuals under older institutions, we mimic this temporal interdependency by having our subjects play all three new versions of our BC game, but with different treatment orders to different subjects without any feedback between games, which indeed brings about interesting differences in out of equilibrium behavior. These order differences constitute a fourth change to the standard BC game.

2 Theoretical Considerations

In the standard beauty contest game with $0 < \beta < 1$ (Moulin (1983)), as first experimentally tested by Nagel (1995), players have to choose a number between zero and 100 inclusive, (allowing for two decimals at most), with the winner being the person(s) whose choice is closest in absolute value to the target of $\beta\times$ the average of all numbers chosen, where $\beta = 1/2, 2/3$, and $4/3$.[7] For the first two cases, the unique, dominance solvable equilibrium is for all persons to guess 0, applying the iterated elimination of weakly dominated strategies. With $\beta = 4/3$, all choosing zero is also an equilibrium, but the only dominant solvable equilibrium is 100. In the case of a tournament payoff structure, where the player(s) who are closest in absolute value to the target value win a large payoff and all others earn nothing as in Nagel (1995) (and most subsequent BC experiments), all strategy combinations are Pareto optimal.

When payoffs are awarded according to a continuous function of the distance between a person's guess and the target number, the dominance solvable equilibrium (by iterative elimination of strictly dominated strategies) in the BC game is the *unique* Pareto optimal strategy.[8] Given a unique Pareto optimal prediction, we can perform a clear-cut welfare analysis across all of our different treatments, which has not been done in previous BC experiments.

Thus far, no experimental beauty contest design has been able to eliminate level-k reasoning immediately, starting from the very first period, when β is different from zero. Allowing subjects to

[7] Most of the experimental literature calls the β parameter p instead of β; we choose to use β to avoid confusion with p-values for statistical significance, as we consider several different value for the β parameter.

[8] Morris and Shin (2002) show that adding a constant to the target number, which is known only approximately with a private and public signal produces a non-Pareto optimal equilibrium, where the inefficiency depends on the noise of the signals.

choose numbers over an unbounded interval, as is allowed in many theoretical models and which we permit in our new version of the BC game, removes the method of iterated elimination of dominated strategies or iterated best reply strategies to find an equilibrium. In particular, subjects no longer have a common anchor, the upper bound of the interval or the midpoint of the interval outside of equilibrium. However, a choice of 0 might become the common anchor. Level 0 types might now propose choosing zero as a naive choice or as the average choice. As a consequence, all subsequent level-k strategies would then also collapse to zero allowing for some noise, of course. Clear spikes or modalities at certain numbers other than zero are less likely to be observed in the BC game with an unbounded interval.

The most important treatment change is our introduction of signals as taken from Benhabib et al. (2015). It is well-known in the psychological literature that individuals frequently use anchors or reference points to make choices. These anchors, whether reasonable or not, can be given exogenously, e.g., by the experimenter or a program, or created by the subject in their own mind. In the original BC game, the anchor is typically the midpoint of the bounded guessing interval, and this value often serves as the guess of a Level 0 type in level-k analysis. If a known constant is added to the BC payoff objective function, (as was first done in Gueth et al. 2002) then this constant can serve as a new anchor, however this type of anchor is not typically an equilibrium choice. Therefore, anchors also need to be chosen wisely by the institution or advisor.

In our new correlated signals design, each subject i, prior to forming her guess, receives a signal, e_i, that is an i.i.d. draw from a normal distribution $N(0,\ \sigma^2)$ and this fact is common knowledge. The payoff function in this precise signal case is the difference between the individual's own guess and $e_i + \beta\ \times$ the average of the guesses other than subject i's own guess. Since the mean of all signals is known to be zero, choosing one's own signal, e_i, is an equilibrium strategy. This technique might be called "strategic nudging", or (re)-anchoring. When we allow correlated signals in game theoretic models, we replace the Nash equilibrium by a type of correlated equilibria, interpreting the correlated signals as a kind of advice. When signals are not wisely chosen, they might not serve as a useful coordination devise.

Buehren et al. (2018) show that idiosyncratic signals drawn from a normal distribution with positive constant mean will result in a more complex equilibrium in which signals are not the equilibrium choice. Indeed they show experimentally that behavior is far from equilibrium and level-k reasoning again becomes a good descriptive model.

2.1 The Games

In this paper we report on behavior from new and different versions of the BC game. All three versions have in common the following features. First, there is a fixed group of n players. Second, each player in the group makes a simultaneous choice (or guess), $x_i \in \mathbb{R}$, where the choice of x_i is limited to two decimal points. Third, each player i's objective is to minimize the distance between his own choice, x_i, and a given and known target, x^*. The target, x^*, is always a function of the average choices of the $n-1$ other players, excluding player i, denoted by $\overline{x}_{-i} = \frac{1}{n-1}\sum_{j\neq i}^{n} x_j$. That is, $x^* = f(\overline{x}_{-i})$, where the function f is publicly known to all.[9] Finally, the payoff to player i in points is given by: $\pi_i = 100 - (x_i - x^*)^2$. Points convert into money earnings according to a known

[9]We make the payoff depend on the average choice excluding player i's own choice so as to eliminate any incentives by player i to manipulate the target value that he or she faces.

constant, with higher points yielding higher money earnings.

Details of the three new beauty contest (BC) game treatments we report on in this paper are as follows:

1. **The new, baseline BC game (called BC)**. This is the same as the original BC game but with choices allowed from the set of all real numbers (i.e., there is no bounded guessing interval). The target, $x^* = \beta\overline{x}_{-i}$, where $-1 < \beta < 1$ is as in the original BC game, but note again that the mean excludes player i's own choice. The equilibrium for this game is that all n players choose zero ($x_i = 0$) and this is the unique, Pareto optimal solution.

The next two treatments introduce idiosyncratic, payoff relevant signals. There can be two equilibria in this setting. When signals are precise then there is a certainty equilibrium where every player chooses their own signal. When signals consist of both an idiosyncratic and a common component, there is a signal extraction problem and a multiplicity of possible equilibria (sunspot equilibria) ex-ante. While in Benhabib et al (2015) the common component is endogenously determined we use an exogenous sentiment shock, for ease of experimental design. As in nature, the realization of the signal with or without a common component is determined before play starts and is known to subjects. More precisely, our two new treatments are as follows:

2. **The BC game with precise signals (called BC-e):**

   In this treatment, the target, $x^* = e_i + \beta\overline{x}_{-i}$, where each individual i's private signal, e_i, is an i.i.d. draw from $N(0, 10)$. The distribution function for signal draws is public knowledge, and in particular all n players know that the mean of the drawn signals is 0. Each subject is informed about his own private signal, but not the signals of others. As in treatment 1, we again allow $-1 < \beta < 1$ and choices from the set of all real numbers (up to two decimals). In the equilibrium of this version of the BC game, each player's guess is equal to his signal ($x_i = e_i$).

3. **BC game with noisy signals and a signal extraction problem (called BC-h):**

   In our more complicated, imprecise signal version of the BC game with signals, the actual signal that subjects receive is a known convex combination of an idiosyncratic signal e_i drawn from a known normal distribution, $F(0, \sigma_F^2)$, and a common-to-all signal, z, that is drawn from another known normal distribution $G(0, \sigma_G^2)$. However, each subject's payoff function continues to depend on their own unobserved idiosyncratic signal alone, that is, we maintain the same payoff function as in the BC-e treatment with target $x^* = e_i + \beta\bar{x}_{-1}$, but the signal of e_i is imprecisely given by $h_i = (1/3)e_i + (2/3)z$. Thus, subjects in our imprecise signal treatment face a more complicated signal extraction problem in making their guesses (see appendix, case noisy signals and extraction). As in treatment 2, e_i is an i.i.d. draw from $N(0, 10)$. The random variable z is another mean zero shock that is common to all n players.

   Note that as detailed in the theoretical section in the appendix section "signal extraction", there is an issue of simultaneity of signals and choices: the signals s_z are observed by players before they choose x, but they contain information about the average play z. For consistency under REE, the optimal actions of players will in fact have to generate an average action $z = X$ for every realization of z.. This only holds if the variance of z is related to the variance of e, as detailed in the appendix, equation (1).

However, in the experiment we exogenously chose z, i.e., it is given through the hint h as described in the appendix "case with noisy signals and signal extraction." In order to choose the variance of z, we use equation (1) of the appendix although any other variance could also be used. As a result, the equilibrium choice is given by $x_i = \gamma h_i$, were γ is independent of β and $\gamma = 1/(2/3) = 1.5$ for all β. The signal h_i is a so-called "sunspot" signal; subjects must employ signal extraction reasoning to make use of this signal. There also exists another equilibrium in which the players ignore the common-to-all signal, z, and all simply choose zero, since subjects know the expected value of the private signal, e_i, which is an i.i.d. draw from $N(0, 10)$. Thus, our setting with noisy signals and a signal extraction problem yields also an interesting equilibrium selection problem. Indeed, as we will see later, the modal choice of our subjects is to ignore the hint h altogether and to play 0 (see also the discussion in the Appendix).

In Figure A.1. of the Appendix, we show the equilibrium realizations for all treatments. Allowing choices over the entire real line was introduced for mathematical reasons since the signals are drawn from a normal distribution. Furthermore, since we also consider negative values for β (the negative feedback case), negative numbers must also be allowed. Having positive and negative values for β can be given an interpretation in terms of strategic complements or strategic substitutes without any additional design difference (see, e.g., Hommes et al. 2013).

3 Experimental design

Each of the three versions of the game was played a single time (one-shot) under 8 different values for the parameter, β. All subjects played all three versions/treatments and were exposed to all 8 parameters within the same game. Thus, for each subject, we have data on 3×8 or 24 one-shot games in total. There was *no feedback* given after each game to avoid learning effects; subjects were only paid on the basis of their point totals following completing of the final, 24th game. Thus we can treat each individual as an independent observation. We use a so-called block design, which means that all 8 parameters within the same game/treatment are played within the same block, and then the next game starts with the same parameter vector for β. Specifically, for each block/treatment we used the set of parameters $\beta = \{-0.9, -2/3, -0.5, -1/3, 1/3, 0.5, 2/3, 0.9\}$. To avoid inducing some monotonicity of behavior (as might arise from having an order of play where β varied from lowest to highest) we changed the parameter values for β from game to game, however, starting all sequences with $\beta = 2/3$. Specifically, the chosen sequence of parameter values was

$$\beta = \{2/3, 1/3, -0.5, -0.9, -1/3, 0.9, -2/3, 0.5\}$$

and we kept this order constant across all three treatments and sessions of our experiment. The subjects learn in block 1 (called stage 1) in each period which β parameter is implemented for that period (game) but not prior to that period (game).

Furthermore, as noted earlier, we introduce four new features to the original BC design:

1. In all three treatments, the choice interval is unbounded; any number in $\mathbb{R}$ up to two decimal points is allowed. Specifically, subjects were instructed that they can "choose any number,

(positive, negative, and zero)." Relative to the original BC game this change serves to de-anchor beliefs.

2. A quadratic loss payoff function, given by: $\pi_i = 100 - (x_i - x^*)^2$, where the target x^* value is either $\beta\overline{x}_{-i}$, as in the BC treatment or $e_i + \beta\overline{x}_{-i}$, as in the BC-e and BC-h treatments. Recall that the average, $\overline{x}_{-i}$, always excludes the individual's own choice, x_i. In the instructions (see the appendix) we also add a table, showing a few possible payoff outcomes (in points) depending on the quadratic distance between x_i and the target for the treatment, x^*.

   Subjects were instructed that one choice (round) would be randomly chosen from each of the three stages (treatments) for payment. As they did not know which round would be chosen in advance, they were incentivized to make choices as close to the target in every round of all three treatments (stages).

3. In treatment 2, a precise idiosyncratic payoff relevant signal, e_i, is introduced to re-anchor beliefs. Subjects are instructed that their signal and every other player's signal is a random, i.i.d. draw from a normal distribution, $N(0, 10)$. To make the signal distribution clear to subjects we presented them with a graph shown in Figure 1 (and taken from the instructions) of this distribution as well as a characterization of the meaning of this distribution in probabilistic terms to the right of the graph.

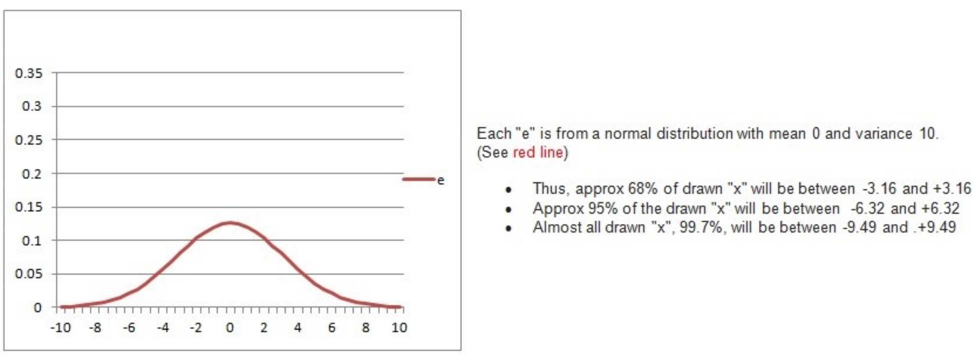

Figure 1: Normal distribution $N(0, 10)$ for the BC-e treatment as shown in the instructions.

4. In treatment 3, the payoff relevant signal, e_i, is imprecise. Subjects were instead given a signal, h_i, and instructed that "You will not know the precise value of e. Instead you receive a hint, h, about the value of e which is specific to you." Subjects were further instructed that the value of h for player i is determined by the formula $h_i = 1/3(e_i) + (2/3)z$, and that z was a common signal for all n group members that was randomly chosen from a different distribution. Specifically, subjects were told that "The value of z is drawn from a normal distribution with mean 0 and a variance that changes from decision to decision." The precise variance of the z variable each round was made clear to subjects prior to their making a choice. The distribution of both e and z were presented to subjects in a graphical form as in the Figure 2 with treatment 3, we explore the extent to which subjects can do the necessary signal extraction needed to arrive at a best guess for their own private signal, e_i.

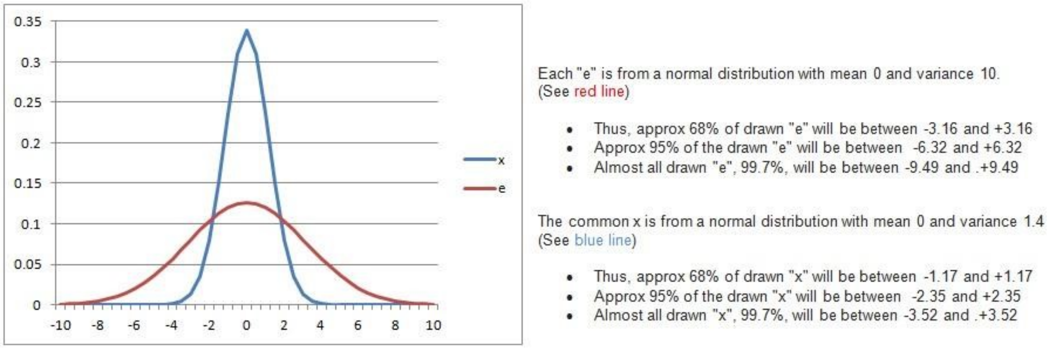

Figure 2: Normal distributions for signal e(i) and the common signal z, as seen in the instructions of treatment BC+h.

Subjects are exposed to all constructed games (called stages), but we vary the order of games in different sessions, using the so-called Latin Square, that is each rule is played once in each position, resulting in 3 different orders. Specifically we have:

Treatment order 1: BC; BC-e; BC-h; with 8 different parameters for β within each stage; 40 subjects.

Treatment order 2: BC-e; BC-h; BC; with 8 different parameters for β within each stage; 40 subjects.

Treatment order 3: BC-h; BC; BC-e; with 8 different parameters for β within each stage; 40 subjects.

For the BC-h treatment we pre-drew, for each β-parameter, two different z values. Twenty (one-half of) subjects in the BC-h treatment faced one of these z-values and the other twenty faced the second z-value.

We conducted two independent sessions for each version of the game (BC, BC-e and BC-h) with 20 subjects in each session. Since there was no feedback between any round, each subject's choices constitute an independent observation from the choices of other subjects. We thus have 40 independent observations (consisting of 3×8 observations by each subject of each treatment. Choices by each subject across treatments are, of course, not independent.

Sessions were conducted at the Pittsburgh Experimental Economics Laboratory with subjects entering their choices using networked computer workstations. The program used was Qualtrics, an online platform typically implemented for survey questions. Since we did not give any feedback to subjects between the different periods, interactive game software, such as zTree, was not necessary.[10]

[10] A further advantage of Qualtrics is that it can be played online by audiences before talks, for teaching purposes, and in demonstration experiments.

4 Experimental findings

Our experiment has yielded a number of interesting findings which we report as a series of different results. Our analysis is guided by the theoretical considerations presented in section 2. Among other predictions, the mean guess in all treatments aside from BC-h should be zero, and the variance of guesses in BC-e and BC-h should be non-zero. Finally, there should be no treatment order effects.

Result 1 *Across all treatments and stages only 1.4 % of choices exceed* $|100|$.

In the remainder of the paper, we will exclude such choices from our analysis as we regard these choices to be extreme outliers.[11].

4.1 First stage choices

While there was no feedback on outcomes, we nevertheless observe strong order effects from exposure to the different treatments, i.e., depending on which of the three games (BC, BC-e, BC-h) is played in the first stage.[12] For this reason, we cannot pool the data from the same game played in different stages. We therefore first report results only from the first stage (8 rounds) of a given treatment order (1, 2 or 3) and compare choices made in the different games played in the first stages of each order. We also make comparisons with choices made in the original BC game with a positive, bounded guessing interval [0,100] and $\beta = 2/3$ using data taken from Nagel (1995). Our new unbounded guessing interval treatment begins with $\beta = 2/3$ and therefore initial conditions are similar to the bounded interval game by Nagel (1995).[13] In what follows, Nagel's (1995) data is referred to as treatment order 0; the data from that study is limited to the first round only and for the $\beta = 2/3$ treatment, which we compare with the corresponding one-shot version from our BC game (order 1, stage 1) with an unbounded guessing interval and also with $\beta = 2/3$.

4.1.1 Removal of bounded choice interval

Consider first the behavior in the BC game of our new environment without signals and without the [0,100] boundary. Using the data from the first stage of our treatment order 1 (where the BC game is played first) we have the following:

Result 2 *Choices in the* $\beta = 2/3$ *game are significantly higher (p-value* < 0.01*) in the original BC game with a closed guessing interval of* $[0, 100]$ *than when the choice interval is unbounded and comprised of the entire real line.*

[11] Excluding those outliers is done to improve graphical exposition, but does not change basic qualitative results below

[12] We address order effects in section 4.1.3.

[13] We do not compare results from bounded $[0, 100]$ interval BC games involving $\beta \neq 2/3$ since first, our new sessions all have $\beta = 2/3$ in the first round and we wanted to compare first round behavior with and without a bounded guessing intervals and second, a guessing interval of $[0, 100]$ may no longer be appropriate in that case where $\beta < 0$ as it would prevent subjects from guessing negative numbers.

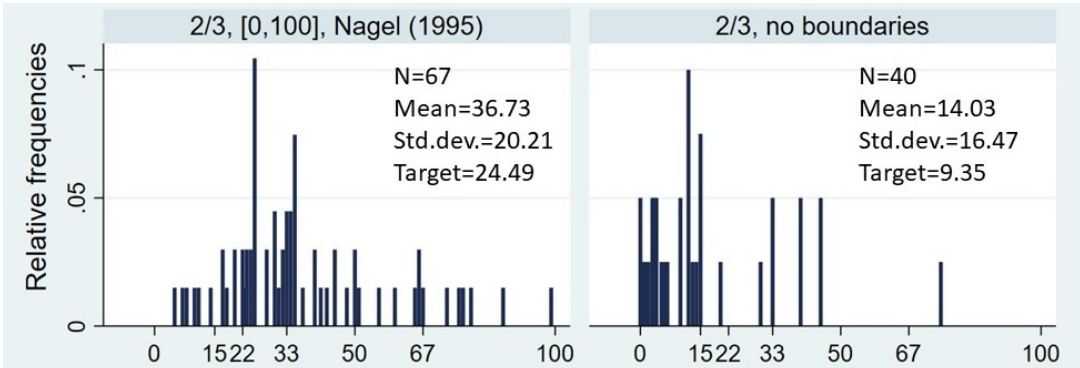

Figure 3: Relative frequency of choices when $\beta = 2/3$: Left panel with boundaries [0,100] from Nagel (1995); right panel without boundaries and no signals.

Support for Result **??** comes from Figure 3 which shows the relative frequencies of choices, the mean, target, and the standard deviation, using the data of the first round of the $\beta = 2/3$ guessing game with a bounded choice interval of $[0, 100]$ taken from Nagel (1995) (referred to as order 0 later on) with data from our first new treatment with an unbounded interval, treatment order 1, where the BC game without signals is played first. Comparing the distribution of guesses for the two versions of the BC game where $\beta = 2/3$, the distributions under the bounded interval are significantly higher than for the unbounded interval according to a Kolmogorov-Smirnoff test (p-value $< .01$).

Result 2 is particularly interesting since the theoretical solution in both games (unbounded or bounded intervals) is the same, namely that all choose zero. In the closed interval experiments, the midpoint of the guessing interval is 50, serving as an intuitive anchor. Since subjects in that environment use low levels of reasoning (typically between levels 0 and 3), choices are far from the equilibrium. Less than 10% choose below or equal to 10.

By contrast, in the unbounded case, the focal point is zero (the midpoint of all possible choices) and thus, in theory all level-k types coincide with choosing zero as well. Consequently, the equilibrium prediction that all choose 0 becomes more focal. Yet, we see still see some deviation from zero, but only towards positive numbers, since the target of $\beta = 2/3$ times the average suggests that choices will be non-negative. Indeed, we observe that 55% of choices in the 2/3 BC game with unbounded interval (order 1) lie in the 0-10 interval and there are no negative choices in this $\beta = 2/3$ treatment (see also appendix Figure A.2.). Thus, our new treatment, with an unbounded interval results in choices that are closer, but not all the way to equilibrium with means around 14. We next turn to analyzing behavior in the two new signal treatments.

4.1.2 Adding idiosyncratic signals

In the analysis that follows, we will no longer use BC choices with a closed interval for comparison purposes, except in section 4.1.4 on welfare analysis. All remaining comparisons will involve BC choices under an *un*bounded guessing interval from the three experimental treatments conducted for this paper.

We present results from the BC, BC-e and BC-h games when each game is played in the *first* stage of an experimental session (order 1, 2 and 3, respectively). Thus, we consider the three games, BC, BE-e and BC-h when subjects had no prior exposure to any other game form in the session. In the remainder of the paper we will graphically present our data with box-plots instead of relative frequency distributions, since only important data variables are the medians, quartiles, variances or outliers.[14]

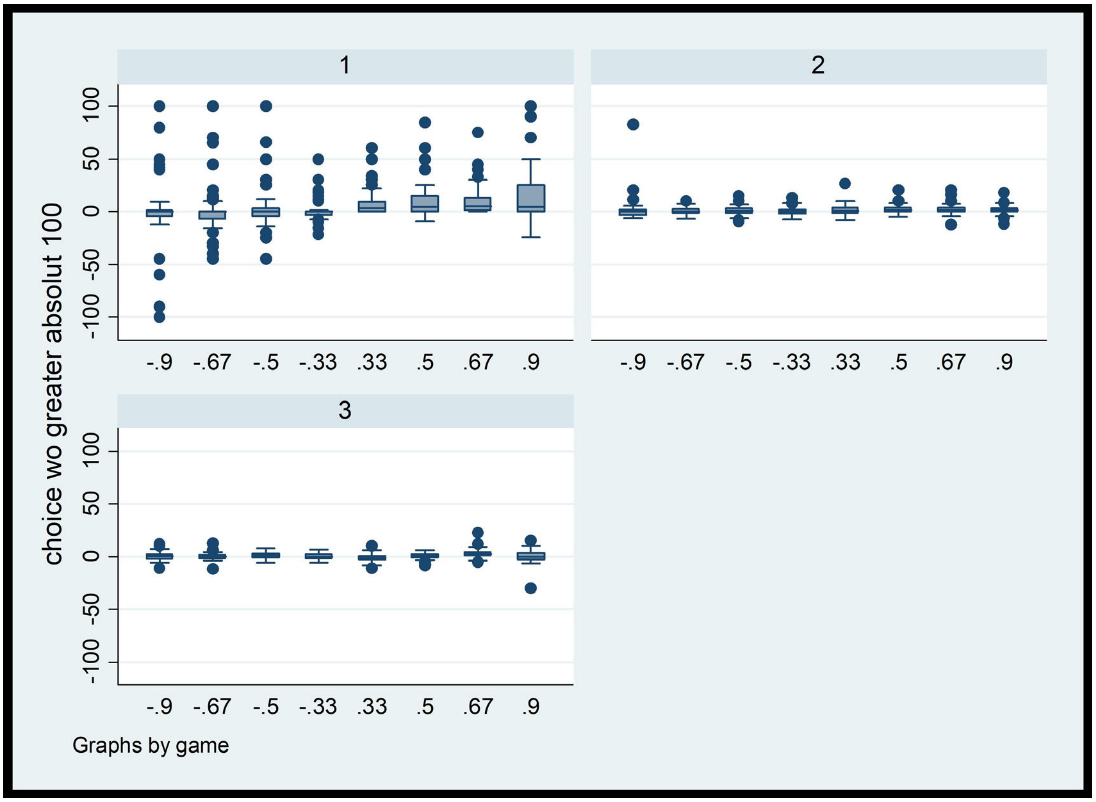

Figure 4: Boxplots for the first stage choices for each β-parameter (-0.9,...0.9); BC (order 1, stage 1), BC-e (order 2, stage 1), BC-h (order 3 stage 3). Choices $> |100|$ are eliminated. There are many more outliers in order 1 as compared to orders 2 and 3.

Result 3 *Regarding choices in the BC, BC-e and BC-h games when played in the first stage:*

1. *The variance in choices is greater in the BC game than in the BC-e or BC-h games which is*

[14]As a reminder, a box-plot provides a five-number summary of a distribution, which in our case amounts to the numbers chosen excluding those numbers that exceed 100 in absolute value. For each box-plot, the lowest horizontal bar represents the lowest number chosen, the next highest horizontal bar, (which is also the lower bound of the box), represents upper bound to the first quartile of chosen numbers, the next horizontal bar is the median number choice (also the upper bound to the second quartile of number choices), the upper bound of the box represents the upper bound to the third quartile of numbers chosen and the final horizontal bar represents the upper bound to the fourth quartile of number choices. The black dots are suspected "outliers" (one per black dot) which are defined as being in excess of 3 times the interquartile range of variation (or third percentile), i.e., the box between the first and third quartiles.

contrary to the theoretical prediction of zero variance in the BC game.

2. *Median choices in the BC game are increasing in $\beta > 0$ but this is not the case in the BC-e or BC-h games.*

3. *For the BC game, when $\beta > 0$ (all orders), 95.6% of choices are positive.*

4. *For the BC game, when $\beta < 0$ (all orders), 44.7% of choices are positive.*

Support for result 3.1 comes from Figure 4 which shows the distribution of choices in box-plots for each β-parameter within each of the three different games, BC, BC-e, BC-h when played in the first stage. While all BC-games without signals contain a lot of outliers, the games with (un)precise signals contain almost none and thus signals are effective at decreasing the distance to equilibrium even more than the feature of having an unbounded choice interval. All single tests of the variances of BC against BC-e or BC-h produce p-values < 0.000. There are no systematic patterns when testing BC-e against BC-h, meaning that some tests of variances are significantly higher, lower, or not significant.

Support for result 3.2 comes from a regression of choices on the different game values for β-parameter> 0 for the BC game which yields a positive slope coefficient (p-value$= 0.012$) and insignificant constant. By contrast, for the other two games, BC-e and BC-h, the slope coefficient is insignificantly different from zero (p-value $> .10$).Furthermore, Mann-Whitney U-tests confirm that for positive β-parameters, BC stage 1 choices are significantly higher than choices of BC-e, stage 1 (order 2) or BC-h choices stage 1 (order 3) with a p-value of 0.0001; For negative parameters, choices are also significantly different between BC and BC-e or BC-h, but with lower significance (p-values < 0.0208). There are no clear differences between BC-e and BC-h choices.

Support for results 3.3 and 3.4 comes from level-k reasoning. As stated in the theoretical section, it is well known (e.g. from Hommes et al. 2013) that in games of strategic substitutes (as in our BC games with $\beta < 0$), behavior is closer to the equilibrium or converges faster to the equilibrium than in games with strategic complements (BC games with $\beta > 0$). Level-k reasoning can explain this phenomenon: If $\beta < 0$ and a subject expects the average will be negative, then her best reply will be a positive choice and vice versa for a positive expectation for the average.[15] With two types, level 0 and level 1, present in subjects' reasoning, therefore, the average is close to zero. This is confirmed by observing 45% positive choices for $\beta < 0$ while for $\beta > 0$ most choices are positive. In games with strategic complements behavior ($\beta > 0$ naive behavior suggests positive numbers and the feedback will lead to higher choices, if one beliefs that other choose higher numbers.

4.1.3 Treatment order effects

We next consider choices in the three different *stages* of each treatment order.[16] The Figure 5 shows box-plots of choices, separately for the three stages of each treatment order, 1,2 or 3, pooled

[15]Whether subjects start with negative or positive guesses for the mean value in their heads and then iterate can best be tested by either eye-tracking data or other cognitive process procedures. We hypothesize that they start with negative numbers as the negative β-parameter suggests.

[16]We do not find evidence for order effects of the different β parameters within a treatment, as was also the case in Coricelli and Nagel (2009).

over all β-parameters within the same game (BC, BC-e, BC-h) (3 games $\times 3$ stages $\times 3$ orders). In the appendix figures A.2-4, we show disaggregated box-plots for each β-parameter in each stage and order.

A particularly striking feature of figure 5 is that the variances of choices in the subsequent stages in order 1 (see three box-plots of left part of figure 5) are much larger as opposed to the other two orders (middle part and right part) when signal games are played first, a finding we summarize as follows:

Result 4 *1. When comparing choices of the same game (BC, BC-e, or BC-h) and parameters across the different orders (e.g. BC-e in order 1 with BC-e in order 2), means (for positive β), variances, and the number of outliers are lower, when signal games are played first (as in order 2 or order 3), than when the BC game without signals is played first (order 1).*

Similarly,

2. Means for $\beta > 0$ and variances of BC without signals in order 2 and order 3 are significantly lower than those in BC-e and BC-h in order 1.

Support for Result 4 comes from Figure 5 and from Figures A.2–A.4 for disaggregated β-parameters in the Appendix. All tests of variances show significant differences of BC (order 1) against BC (order 2 or 3); BC-e (order 1) against BC-e (order 2 or order 3); BC-h (order 1) against BC-h (order 2 or order 3), and BC (order 2 or 3) vs BC-e or BC-h (order 1) with p-values < 0.0001, using any conventional test.

Mann-Whitney U-tests of medians show that for all pairwise comparisons of choices in positive β-parameterizations of BC, BC-e, and BC-h against choices in those same games in orders 2 or 3 (e.g., BC, β=2/3, order 1 vs BC, β=2/3 order 2) and also BC order 2 or 3 vs BC-e order 1 are significantly different (p-value < 0.0086).

This is obviously due to the large number of outliers present in all games in order 1. The signals in order 1 (added in stage 2 and 3) do not have a strong effect as in order 2 and 3 where they are added in stage 1 and thus become a strong reference point (see discussion of previous subsection). Once a subject supposes a mean very different from zero in the first stage of order 1 in a BC game without signal, subsequent signals in stage 2 will just distort that belief. Different beliefs in BC games without signals (thus without anchors) produce a larger variance. Comparisons between the same games of order 2 versus order 3 do not show systematic different patterns.

These results are very interesting as it appears that subjects understand that in order 2 or 3, the standard BC game converts into a BC+0 game, where the anchor of zero is common knowledge, once they have been exposed to the BC-e or BC-h games. Thus, even though there is no information from one situation to the next, BC games played after BC-e or BC-h show similar results as games with signals.

Medians for negative β-parameters are not different between different orders and games. This result comes from level k reasoning in negative feedback settings, which brings the mean closer to zero as discussed in the previous subsection.

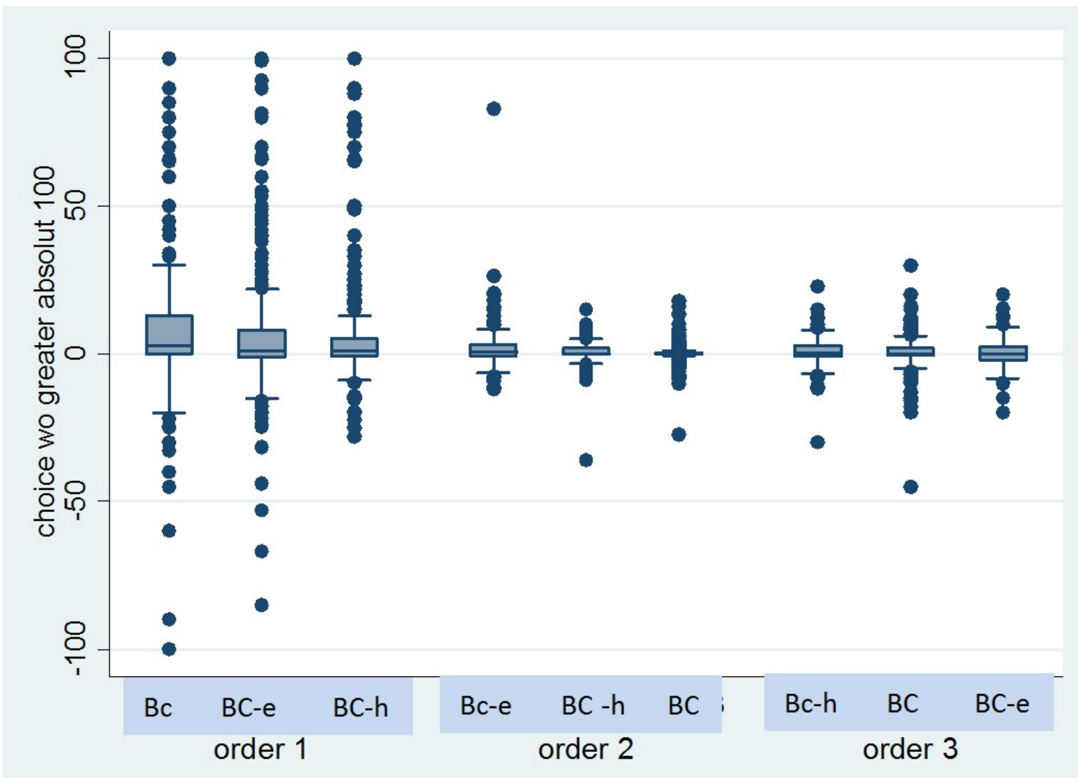

Figure 5: Boxplots of choices of each game for each order, pooling overall β-parameters within each stage and game, respectively. (In the appendix figures A2-4 we show the disaggregated results).

	order 1 e	order 1 h	order 2 e	order 2 h	order 3 e	order 3 h
	stage 2	stage 3	stage 1	stage 2	stage 3	stage 1
mean coefficient	-3.82	2.62	1.16	-1.25	0.87	-3094.60
median coefficient	0.96	0.53	0.63	0.13	1.00	0.70
% $\gamma = 0$	0.05	0.15	0.11	0.27	0.06	0.07
% $\gamma = 1$	0.18	0.02	0.14	0.03	0.14	0.03

Table 1: Mean and median coefficients for γ in the BC-e and BC-h games, and relative frequency of $\gamma = 0$ or 1, separately for each treatment order. The grid width is 0.1 for measuring the γ frequency

These results also indicate that the order of institutional changes matters in the lab. Under order 1, one could argue that idiosyncratic signals do not work very well, when subjects' minds are anchored by the reasoning they applied in the BC game where there are no signals and an unbounded choice interval. However, in subject pools for which such confounding factors have not yet taken place, the institutional implementation of idiosyncratic signals has the desired effect of achieving coordination on the equilibrium where many choose close to their signal or even 0, even in the BC game without such signals. This will also be explored in the next section.

4.1.4 Use of precise and imprecise signals

In this section, we take a closer look at the relationship between the signals and the actual choices in the BC-e and BC-h treatments. As we have seen in the previous sections, order effects matter a lot with respect to the variance in choices and the number of outliers in these games. Therefore, we should expect that signals are used differently in the different treatments. In the BC-e treatment, the equilibrium choice is "choose your signal", e_i. This choice corresponds to an equilibrium coefficient $\gamma = \textit{choice}/\textit{signal} = 1$. In the BC-h treatment with imprecise signals, in one equilibrium, subjects must solve a complicated signal extraction problem resulting in an equilibrium coefficient for $\gamma = \textit{choice}/\textit{signal} = 1.5$. A second, simpler equilibrium requires only that all players ignore the hint altogether and therefore play zero. Note that the equilibrium distribution for γ is degenerate in all equilibria for all games and parameters with precise and imprecise signals.

Figure 6 plots the actual relative frequency distribution of gamma for both games (BC-e and BC-h, separately, given the choice/signal relationship of each subject pooled over all β-parameter constellations. Modal choices are at $\gamma = 1$ in the BC-e game and are at 0 (ignoring signals) in BC-h. There is also a spike at 0 (ignoring the signal) in the BC-e game and a small spike at $\gamma = 3$ in BC-h, which corresponds to players ignoring the *common* part of their signal, z, and inferring from h_i that $e_i = 3h_i$. In the Appendix Figure A.6 we show the different frequency for each order and game, separately.

Result 5 *In the BC-e game, for each order and β value, the modal choice is always equal to the private signal, i.e., $\gamma = 1$. However, in order 1 (where BC-e is played second, after BC) the choice equals the signal only 18% of the time while in orders 2 and 3 (where BC-e is played first or third),*

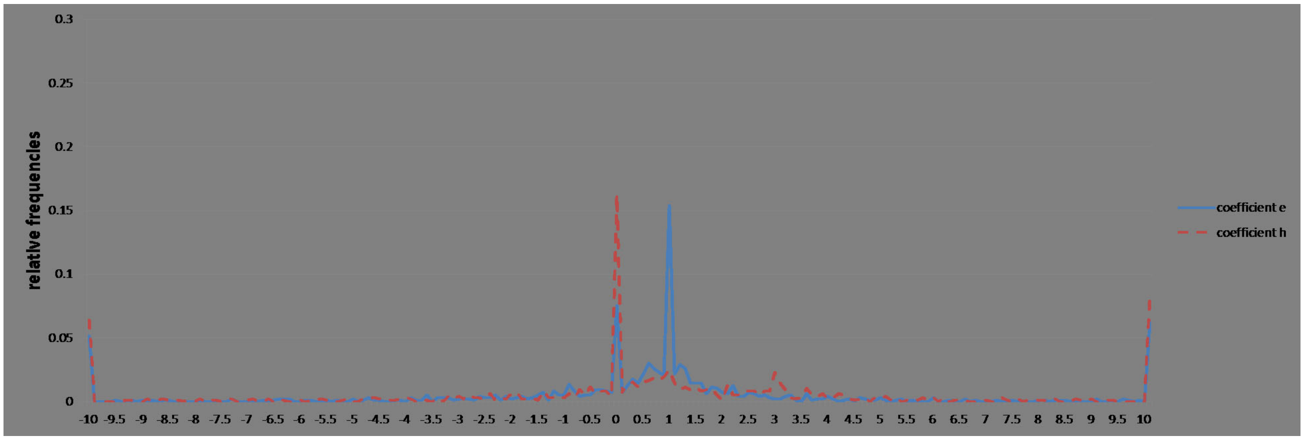

Figure 6: Relative frequencies of γ coefficients (where choice(i)= $\gamma * e(i)$ or $\gamma * h(i)$, respectively). We plot the frequencies for both the BC-e and BC-h-treatments pooled over all stages and β-parameters. (There are in total $3 \times 8 \times 40 = 960$ observations for each of the two lines) In the Appendix we disaggregate the different stages and also add the professional treatment graphs.

the choice equals the signal just 14% of the time. The median γ, however, is generally close to the theoretical prediction of 1. *Further, the tails of the distribution of guesses (coefficient $|\gamma| > 10$) are much larger in order 1 (10% and 15%) than in order 2 or order 3 (less than 5%).*

Support for Result 5 comes from Table 1 which reports the mean/median γ coefficient using the actual choices and the corresponding signals of all subjects, separately for each game and order. If we enlarge the signal compliance interval, to include small "trembles" of $+/-$ 0.4 of the signal, we find that the frequency of $\gamma = 1$ (signal compliance) is 29% of the time in order 1, 38% in order 2 and 41% in order 3. Figure A.6. in the Appendix considers the different order effects which show clear differences in the tails between the different orders which correspond to the different outliers we showed in Figure 5 of the previous section. There is also a spike at $\gamma = 0$ for BC-e, comprising 5%, 11% and 6% of all choices in treatment orders 1, 2 and 3, respectively. The frequency of $\gamma = 0$ choices in BC-e rises to 12%, 21% and 14%, respectively for orders 1, 2, and 3, if we allow for small "trembles" of $+/-$ 0.4 of the signal. This means that there are some systematic choice patterns, but also a high heterogeneity between players within the different treatment orders.

Result 6 *In the BC-h game, the modal choice for the γ coefficient is* 0, *which corresponds to the equilibrium where everyone ignores the signal, accounting for 15%, 27%, 7% of γ's in the three different orders. The resulting median choice is below* 1 *in all three orders and thus is not close to the more complicated signal extraction equilibrium coefficient for γ of 1.5.*

Support for this result again comes from Table 1. If we allow for trembles around $\gamma = 0$ of $+$ 0.4, the frequency of $\gamma = 0$ (ignoring the signal) choices rises to 20% 39% and 15%, respectively for the three orders 1, 2, and 3. A t-test confirms that in all orders, the γ coefficients are significantly different from the 1.5 signal extraction equilibrium predictions ($p > 0.42$, $p > 0.40$, $p > 0.16$ for the three different orders, respectively). Interestingly, there is no spike at $\gamma = 1$ which would mean to choose the hint. There is also some mass (but less than 5%) at a γ coefficient of 3. Such behavior ignores the common part of the signal, z, and, since $h = 1/3e + 2/3z$, yields a choice of $3h$). However, this is not an equilibrium strategy while everyone ignoring their signal, $\gamma = 0$, is an equilibrium.

Thus, in this section we have seen two simple heuristics: 1. choose (close) to your signal in the BC-e-game which accounts for between 29% and 41% of choices (including trembles). 2. ignore the signal (play close to zero) in the BC-e game accounting for between 12% and 21% of choices and in the BC-h game accounting for between 15% and 39% of choices.

4.1.5 Efficiency

As stated in the model section, when subjects are paid according to a distance function of their choice from the desired target, as opposed to having a single fixed prize for the winner, the unique (rationalizable) equilibrium is also Pareto optimal with the highest possible payoff for all subjects. Any deviation from equilibrium reduces the total sum of payoffs. Thus, we can calculate *efficiency gains* or losses across the different games and in the different orders.

For the no-signal treatment, the highest possible payoff is 100 for each player obtained in equilibrium. For the precise signal treatment, the highest payoff a single player can earn is also 100 when his choice is equal to the target number. However, because of the small numbers of subjects in our experiment, it can happen that the average of all chosen signals is not zero. Thus, playing the signal is only correct in expectation, though the deviations are empirically small (less than 1). Thus, we will normalize the highest efficiency to be 100. In the imprecise signal treatment, there is a loss of efficiency also in equilibrium because of the incomplete information about the signal. Given the parameter choices, the efficiency loss is about 2%. For our qualitative results and comparisons we will again ignore this small loss. We report the actual payoffs in box-plots across all subjects in a particular stage and treatment order. These payoffs also serve as realized measures of efficiency, which are maximal at 100 or near 100, when all play equilibrium, the maximal efficiency.

We use the recombinant method (adapted from Abrevaya (2008) to calculate the expected payoffs (the recombinant method is the same for the different payoff types):

1. Generate reference groups: These consist of 10 subjects in the same session, order, game, β, stage and subgroup (subgroup is an auxiliary concept to distinguish subjects 1-10 from 11-20 for computational purposes).

2. For each subject, we create their payoffs for all possible 8-tuples of other subjects within the same reference group.[17]

3. Then, for each payoff type, we take the mean over all possible 8-tuples for each payoff type for each subject.

In Figure 7 we show the actual payoff pooled over all β-parameters for each game within each order, e.g., "BC 2" means: the BC payoff earned in order 2 where BC is played in stage 3. "BC

[17] We use 8-tuples rather than 9-tuples because the computational time complexity of the 8-tuple case is considerably less than in the 9-tuple case, a difference of days versus weeks.

0" are the two games with the bounded [0,100] guessing interval taken from Nagel (1995). "BC 4, BC-e 4, BC-h 4" are results from experiments with 18 economic professors and PhD students before a talk at NYU (by Nagel) using the same software for order 1.

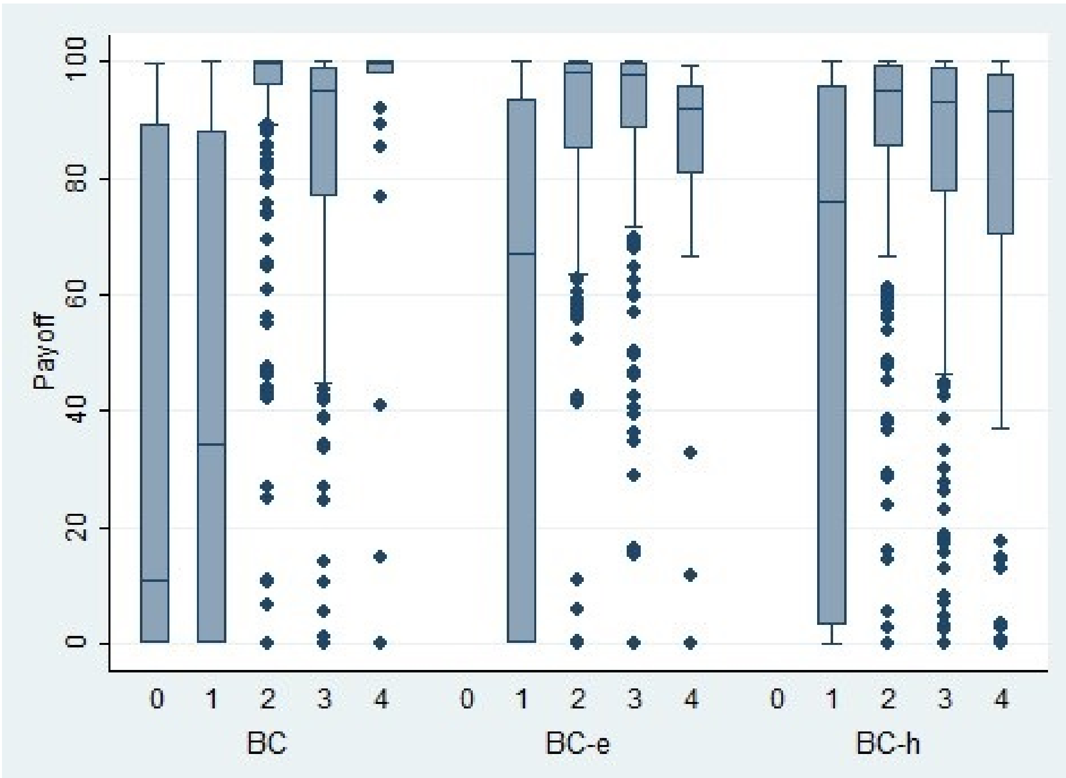

Figure 7: Actual payoffs; the x-axis describes the different orders: 0=Nagel, 1995 data, 1=order 1,.. 4=professionals

Figure A.4 in the Appendix A8-10 shows the same data, disaggregated for each β-parameter, order, and stage; A11 shows the data by Nagel (1995) and A12 the payoffs from 18 economic professors and PhD students before a talk in NYU by one of the authors (Rosemarie Nagel).

From Figure 7, we have the final result

Result 7 *Average efficiency is lowest when there is a bounded guessing interval (Nagel 1995 data). Efficiency is increased when the boundaries are eliminated (order 1, BC), Efficiency exceeds 80% when signals are provided or BC games are played after BC-e or BC-h (order 2 and 3).*

In the original BC game with bounded interval (order 0), and $\beta = 2/3$ efficiency has a mean 35% and median of 12%. When the BC game with an unbounded guessing interval and $\beta = 2/3$ is played in the first stage, efficiency has mean 56% and median 41%. These efficiency differences are statistically significant, p-value=0.006 Introducing signals as in the BC-e and BC-h treatments raises efficiency to over 80% only when those two treatments are played in the first stage (orders 2 and 3).

Pairwise comparisons of the average efficiency payoffs by player in the same games (e.g., BC) across different orders with Kolmogorov-Smirnov tests reject equal distributions in order 1 (BC, BC-e, BC-h) versus order 2 (BC-e, BC-h, BC) (p-value: 0.000) and also rejects equal distributions in order 1 (BC, BC-e, BC-h) against order 3 (BC-h, BC, BC-e) (p-value: 0.000).[18]

5 Conclusion

Understanding the process by which economic agents, with no prior history of interaction, nevertheless coordinate on equilibrium outcomes is important for modeling dynamic adjustments (e.g., to changes in fundamentals) as well as for questions of equilibrium selection. The beauty contest game provides a useful framework for understanding this process as it is a reduced form representation of the coordination problem implicit in many games and market models. In this paper we have made several changes to the original beauty contest game that make it a more general framework for studying this coordination problem. The general finding is that in all of our new treatments, behavior is closer to the unique Pareto optimal equilibrium of the BC game than previously observed in other variations of the BC game, when played for the first time.

The removal of boundaries changes the behavior of subjects in the original BC game with $\beta = 2/3$ and a guessing interval of $[0, 100]$, from a median guess of 36.73 to a median guess of 14.03 in the unbounded version of the game with the same β. The simple reason for this difference is that the midpoint of the bounded interval as a focal point is removed. The equilibrium choice in the BC game of zero now enters as a natural new reference point. The most dramatic change occurs when we introduce precise idiosyncratic and payoff relevant signals, drawn from a distribution with zero mean, which serves as a simple reference point or anchor. An interesting finding is that only about 30% of subjects choose their signals or close to it. Yet, nearly all other mistakes or noisy actions cancel out, since on average, the signal coefficient is close to the theoretical prediction. When signals are imprecise, in one equilibrium, subjects need to perform complicated signal extraction calculations, but this is not what we observe. Instead, subjects fall back to two different kinds of simple heuristics - ignoring the signal altogether and playing zero which indeed is a second equilibrium in the BC-h game. In addition, a few subjects ignore the public signal component. Interestingly, in this last treatment we do not observe subjects using the hint as their choice.

Choosing a signal is an obvious simple heuristic whether it is relevant to the game or equilibrium behavior or not as has been shown in many experiments. Ignoring payoff or equilibrium relevant information seems to be a simple heuristic observed in some experiments. Bayona et al (2016) show that subjects choose (close to) equilibrium actions in supply function games with signals about uncorrelated cost structures between players. The reason is that these equilibrium actions are reached easily through level 1 reasoning procedures. However, players choose similarly also in their correlated costs treatments. They completely ignore that costs and thus signals are correlated. In Ngangoue and Weizsäcker (2015) subjects who receive equilibrium relevant price information given through the behavior of other players in a preceding decision stage ignore that information. Instead they act as their co-players of the preceding stage as if information was not available.

We also find that the treatment order matters substantially. When the original BC game is played first, behavior in the subsequent 'signal' versions of the BC games remains very noisy as

[18] The hypothesis of equal distributions in order 2 (BC, BC-e, BC-h) and order 3 (BC-e, BC-h, BC) is not rejected by Kolmogorov-Smirnov tests (p-value: 0.217).

some subjects seem to consider means that are not (close to) zero and therefore continue to play far away from the equilibrium strategy. These outliers are not present when the BC games with signals (BC-e and BC-h) are played first. Additionally, initial experience playing the signal versions of the BC game has an advantage for later play of BC games without signals, as it makes 0 a strong anchor in those games.

In many models, boundaries on forecast intervals do not exist, or only exist in a single dimension (e.g. prices cannot be negative). Thus, we view our new design as a more natural one in which to study expectation formation. Welfare effects are largely increased in all of our new treatments. Together with our introduction of signals and sentiments, we can reduce and circumvent choices based on the Keynesian level k approach and thus end with Keynes, "calculating where we can, but often falling back for our motive on whim or sentiment or chance." [Keynes (1936), pp. 162-163]. Of course, these sentiments need to be chosen wisely to be payoff maximizing.

References

Abrevaya, J. (2008). On recombinant estimation for experimental data.*Experimental Economics* 11, 25-52.

Alaoui, L., and Penta, A. (2016). Endogenous Depth of Reasoning. *Review of Economic Studies* 83, 1297–1333.

Alaoui, L., and Penta, A. (2017). Cost-Benefit Analysis in Reasoning. mimeo.

Angeletos, G-M., and Lian, C. (2016). Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination. In: J.B. Taylor and H. Uhlig, (eds.) *Handbook of Macroeconomics*, Vol. 2A, Amsterdam: Elsevier.

Angeletos, G-M., and J. La'O, (2012). Sentiments. *Econometrica*, 81, 739-779.

Arifovic, J., Duffy, J., (2018). Heterogeneous Agent Modeling: Experimental Evidence. In: *Handbook of Computational Economics*, vol. 4. Elsevier, pp. 491–540

Baeriswyl, R., and Cornand, C. (2014). Reducing Overreaction to Central Banks' Disclosures: Theory and Experiment. *Journal of the European Economic Association*, 12(4), 1087-1126.

Bayona, A., J. Brandts, and X. Vives, (2016). Information Frictions and Market Power: A Laboratory Study, CESifo Working Paper Series 6005, CESifo Group Munich.

Benhabib, J. (2015). Correlated Equilibria in Beauty Contest Games, mimeo.

Benhabib, J., Wang, P. and Wen, Y. (2015) Sentiments and Aggregate Fluctuations. *Econometrica*, 83, 549-585.

Benhabib, J., Wang, P., and Wen, Y. (2013) Uncertainty and Sentiment-Driven Equilibria. NBER Working Paper No. 18878.

Buehren, C., J. Duffy, and R. Nagel (2018). (Non)-equilibrium anchors in Beauty Contest Games. mimeo

Bergemann, s and S. Morris (2016). Bayes Correlated Equilibrium and the Comparison of Information Structures in Games. *Theoretical Economics*, 11, 487-522.

Bosch-Domènech , A., Montalvo, J.G., Nagel, R., and Satorra, A. (2002). One, two, (three), infinity,...: Newspaper and lab beauty-contest experiments. *American Economic Review*, 92(5), 1687-1701.

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton: Princeton University Press.

Camerer, C. F., Ho, T. H., and Chong, J.K. (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics*, 119(3), 861-898.

Coricelli, G., and Nagel, R. (2009). Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, 106(23), 9163-9168.

Crawford, V.P., Costa-Gomes, M.A., and Iriberri, N. (2013). Structural models of non-equilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1), 5-62.

Farhi, E., and Werning, G. (2016). The Power of Forward Guidance with Bounded Rationality and Incomplete Markets. mimeo

Fehr, D., Heinemann, F., Llorente-Saguer, A. (2019). The Power of Sunspots: An Experimental Analysis. *Journal of Monetary Economics* 103, 123-36.

García-Schmidt, M. and Woodford, M. (2015). Are Low Interest Rates Deflationary? A Paradox of Perfect-foresight Analysis. NBER Working Paper No. 21614.

Güth, W., Kocher, M., and Sutter, M. (2002). Experimental 'beauty contests' with homogeneous and heterogeneous players and with interior and boundary equilibria. *Economics Letters*, 74(2), 219-228.

Hommes, C., and Zhu, M. (2014). Behavioral Learning Equilibria. *Journal of Economic Theory*, 150, 778-814.

Hommes, C., Sonnemans, J., Tuinstra, J., and Van de Velden, H. (2005). Coordination of Expectations in Asset Pricing Experiments. *Review of Financial Studies*, 18(3), 955-980.

Keynes, J.M. (1936). *The General Theory of Interest, Employment and Money*, London: Macmillan.

Li Shengwu, (2017). Obvious Ex Post Equilibrium. *American Economic Review*

107(5), 230-234.

Ledoux, A. (1981). Concours résultats complets. Les victimes se sont plu à jouer le 14 d'atout, *Jeux et Stratégie* 2(10), 10-11

Lucas. R.E. Jr. (1972). Expectations and the Neutrality of Money *Journal of Economic Theory* 4(2), 103–124.

Morris, S., and Shin, H.S. (2002). Social value of public information. *American Economic Review*, 92(5), 1521-1534.

Moulin, H. (1979). Dominance solvable voting schemes. *Econometrica*, 47(6), 1337-1351.

Moulin, H.(1986). *Game Theory for Social Sciences.* New York: New York University Press.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5), 1313-1326.

Nagel, R., Bühren, C., and Frank, B. (2017). Inspired and inspiring: Hervé Moulin and the discovery of the beauty contest game. *Mathematical Social Sciences* 90, 191-207

Ngangoue, K. and G. Weizsäcker. (2015). Learning from Unrealized versus Realized Prices, Discussion Papers of DIW Berlin 1487, DIW Berlin, German Institute for Economic Research.

Nimark, K.P. (2011). Dynamic Higher Order Expectations, mimeo.

Slonim, R.L. (2005). Competing Against Experienced and Inexperienced Players. *Experimental Economics*, 8(1), 55-75.

Stahl, D.O., and Wilson, P.W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1), 218-254.

Sutan, A. and Willinger, M. (2009). Guessing with Negative Feedback: An Experiment. *Journal of Economic Dynamics and Control*, 33, 1123-1133.

Thaler, R. and C. Sunstein (2008). *Nudge: Improving Decisions About Health, Wealth and Happiness*. New York: Penguin Books.

Tversky, A. and D. Kahneman (1986). Rational Choice and the Framing of Decisions. *Journal of Business* 59(4) Part 2, S251-S278.

Appendix

The Game

This appendix departs from the working paper Benhabib (2015).

There are n players is indexed by i, where we assume n is a large number. Each player chooses an action $x(i) \in (-\infty, \infty)$. The average action of all players is represented by $X \in R$, where $X = \frac{1}{n}\sum_{1}^{n} x(i)$ and $x = (x(1), x(2) ... x(n))$. The payoff to player i who's choice is $x(i) \in (-\infty, \infty)$ decreases with the distance $|x(i) - (\beta * X + e(i))|$, where β can take any value and $e(i)$ is a random variable, $e(i) \sim G(c, \sigma_e^2)$, $c \geq 0$. We assume $G(c, \sigma_e^2)$ is common knowledge. Before making their choice players receive signals (sentiments or sunspots), $s_z(i) = az + be(i), 1 = 1...n$, where a and b are non-negative constants and where it is common knowledge that z is a non-fundamental extrinsic random variable, $z \sim F(0, \sigma_z^2)$. We define a signal profile as $s_z = \{s_z(1), s_z(2), ... s_z(n)\}$. In equilibrium each player believes that z will correspond to the average action X, that is that $z = X$. Note that the signals s_z are observed by players before they choose x, but they contain information about the average play z. For consistency under REE, the optimal actions of players will in fact have to generate an average action $z = X$ for every realization of z. It also follows that if the game is repeated over time the average action X will have the distribution $F(0, \sigma_z^2)$, confirming the beliefs of players. As we show later, this will be possible only for equilibrium beliefs about z, that is only for some particular σ_z^2 defining the distribution of $z \sim F(0, \sigma_x^2)$.[19] In our experiment we will use a version if imprecise signals in which the common component is exogenenously given as stated in the last subsection at the end of this appendix.

The information structure of this BC game differs from those recently used in examples of linear-quadratic games with imperfect information. In these games payoffs to players depend on actions and on a true fundamental state of the world, say θ with distribution φ, as for example in Bergemann and Morris (2013). Bergemann and Morris (2013) define a Bayesian Correlated Equilibrium as a distribution μ over (x, X, θ) such that i) players' choices are best responses, and ii) the consistency requirement, that the marginal distribution $\mu_\theta = \varphi$, is satisfied. In our context the random variable z that also enters the signal profile s_z is not a realization of the true state of the world, but it is the endogenous average choice. Thus our consistency requirement under rational expectations is $z = X$ for every realization of z drawn from the distribution $F(0, \sigma_z^2)$ which is common knowledge.

The optimal action for player i in our game is $x(i) = E(\beta * X + e(i)) | s(i))$. The number players, n, will be large enough so we can assume $n^{-1}\sum_{n} e(i) = c$.

Correlated Equilibria

Revealing Signals

We first consider the simplest case is the BC game where sentiments z play no role. Assume $e(i) \sim G(c, \sigma_e^2) \geq 0$, $c \geq 0$. Players are given a revealing signal $s(i) = e(i)$,, which in the

[19] For related context and results see Angeletos and La'O (2013) and Benhabib, Wang and Wen (2015). In some cases as discussed in Benhabib, Wang and Wen (2013) there are a continuum of such sentiment-driven rational expectations equilibria, parameterized by the variance of sentiments.

experiment is called perfect signal, or they ignore z, and they focus only their idiosyncratic shock $e(i)$. The signal fully reveals $e(i)$, and the correlation device z plays no role. With large numbers, that is large n, the players correctly expect $\frac{\sum_n e(i)}{n} = c$. Then it is easy to show, if $0 \leq \beta < 1$, that equilibrium play is

$$x(i) = \frac{\beta c}{1-\beta} + e(i)$$

and

$$average\ play = X = \frac{1}{n}\sum_i x(i) = \frac{1}{n}\sum_i \left(\frac{\beta c}{1-\beta} + e(i)\right) = \frac{c}{1-\beta}$$

consistent with private signals $e(i) \sim G\left(c, \sigma_e^2\right)$. If however we further restrict this game so $\beta > 1$ and $x(i) \geq 0$, the above equilibrium is ruled out since average play $X = \frac{c}{1-\beta} < 0$.

Signal Extraction

Let's now consider the original game with $x(i) \in (-\infty, \infty)$, where $z \sim N\left(0, \sigma_z^2\right)$ and $e(i) \sim N\left(0, \sigma_e^2\right)$, where we set $c = 0$, and where players get a signal $s(i) = az + be(i)$,with $a, b > 0$. In a rational expectations average choice X will have to equal z for every realization of z. If for large n, we have $\sum_{j \neq i} e(j) \sim 0$), the best response $x(i)$ of player i, assuming for the moment that everyone else reports

$$y(j) = \gamma s(j)$$

is given by:

$$\begin{aligned}
x(i) &= \frac{cov\left(\beta * average\ play\ of\ others + e(i), s(i)\right)}{var(s(i))} \\
x(i) &= \frac{cov\left(\frac{1}{n-1}\beta \sum_{j \neq i}(\gamma s(j)) + e(i), s(i)\right)}{var(s(i))} \\
x(i) &= \frac{cov\left(\frac{1}{n-1}\beta \sum_{j \neq i}\gamma(az + be(j)) + e(i), s(i)\right)}{var(s(i))} \\
x(i) &= \frac{\left(\beta\gamma a^2\sigma_z^2 + b\sigma_e^2\right)}{a^2\sigma_z^2 + b^2\sigma_e^2} s(i)
\end{aligned}$$

So in equilibrium we must have

$$\frac{\left(\beta\gamma a^2\sigma_z^2 + b\sigma_e^2\right)}{a^2\sigma_z^2 + b^2\sigma_e^2} = \gamma$$

which simplifies to

$$\gamma = \left(a^2\sigma_z^2(1-\beta) + b^2\sigma_e^2\right)^{-1} b\sigma_e^2,$$

so players' choices, γs_z, depend on the realization of z. Therefore we have a continuum of equilibria with different average actions X. For all j the optimal choice $\gamma s(j)$ minimizes $|x(j) - (\beta * z + e(j))|$. Remember that prior to making his choice $x(j)$, player j does not see z or $e(j)$, but sees only $s(j)$. Now we can derive the restriction on σ_z^2, or on beliefs, so that the average action X is consistent with the realization of the sentiment z, that is $X = z$.

The distribution of z for which players obtain a signal will now be treated as the belief that induces the expectations to be correct. In other words we need to find a restriction on the beliefs for z, parameterized by σ_z^2, for which players are always right, irrespective of the realization of z,

$$z = \frac{\sum_n \gamma s(i)}{n} = \frac{\sum_n \gamma (az + be(i))}{n} = \gamma a z.$$

The last equality follows if there are a large number of players so in the limit as n gets large, $\sum_n e(i) = 0$. We maintain the assumption for the time being that $x(i) \in (-\infty, \infty)$. The above can only be true if $\gamma a = 1$:

$$a\left(a^2\sigma_z^2(1-\beta) + b^2\sigma_e^2\right)^{-1} b\sigma_e^2 = 1$$

$$\begin{aligned} ab\sigma_e^2 &= \left(a^2\sigma_z^2(1-\beta) + b^2\sigma_e^2\right) \\ (a-b)\, b\sigma_e^2 &= \left(a^2\sigma_z^2(1-\beta)\right) \end{aligned}$$

If $(a-b)(1-\beta) > 0$,

$$\sigma_z^2 = \frac{(a-b)\, b\sigma_e^2}{a^2(1-\beta)} > 0 \tag{1}$$

$$\gamma = \left(a^2\frac{(a-b)\, b\sigma_e^2}{a^2(1-\beta)}(1-\beta) + b^2\sigma_e^2\right)^{-1} b\sigma_e^2$$

or

$$\gamma = a^{-1}$$

Therefore, if we want to satisfy REE so that the average play X is equal to z for each realization of z we need equation 1 to hold. Since $b > 0$, we have $\sigma_z^2 > 0$ if either $\beta < 1$ and $a - b > 0$, or if $\beta > 1$ and $a - b < 0$.

Case with noisy signals and signal extraction

Assume again that $c = 0$, and that the agents receive the signal $s(j)$ as above, but that z is an exogenous noise, not related to sentiments about average play. In our experiment we use this specification. It is no longer true that "In equilibrium each player believes that z will correspond to the average action X, that is that $z = X$." The agent j has the objective function to minimize $|x(j) - (\beta * X + e(j))|$, trying to match β times average action X plus his idiosyncratic shock, and will still optimally play

$$y(j) = \gamma s(j) = \gamma(az + be(j))$$

since he does not observe z or $e(j)$. But now there is no equilibrium requirement,

$$z = \frac{\sum_n \gamma s(i)}{n} = \frac{\sum_n \gamma (az + be(i))}{n} = \gamma a z \text{ and } \gamma a = 1$$

In effect, z is just a confounding factor for agent j who cannot directly observe $e(j)$. So we do not restrict the variance of z, σ_z^2, as we did before. [20]

If the variance is unrestricted, the average action will differ from the forecast of agents depending on the realization of their own signal. But z is not the self-fulfilling average action now, and unlike before, $X \neq z$ each period even though $e(j)$ has zero mean. We no longer have to refer to self-fulfilling beliefs about average action given the restriction on σ_z^2. So we are outside the strict BC framework of Benhabib et al. (2015), but we still have a signal extraction and RE. There is also another equilibrium in which the agents ignore the signal z and since the expected mean of e is zero all play zero.

[20] Note that in our experiment we specify the variance of z as in the section of signal extraction with sentiments.

Additional Figures

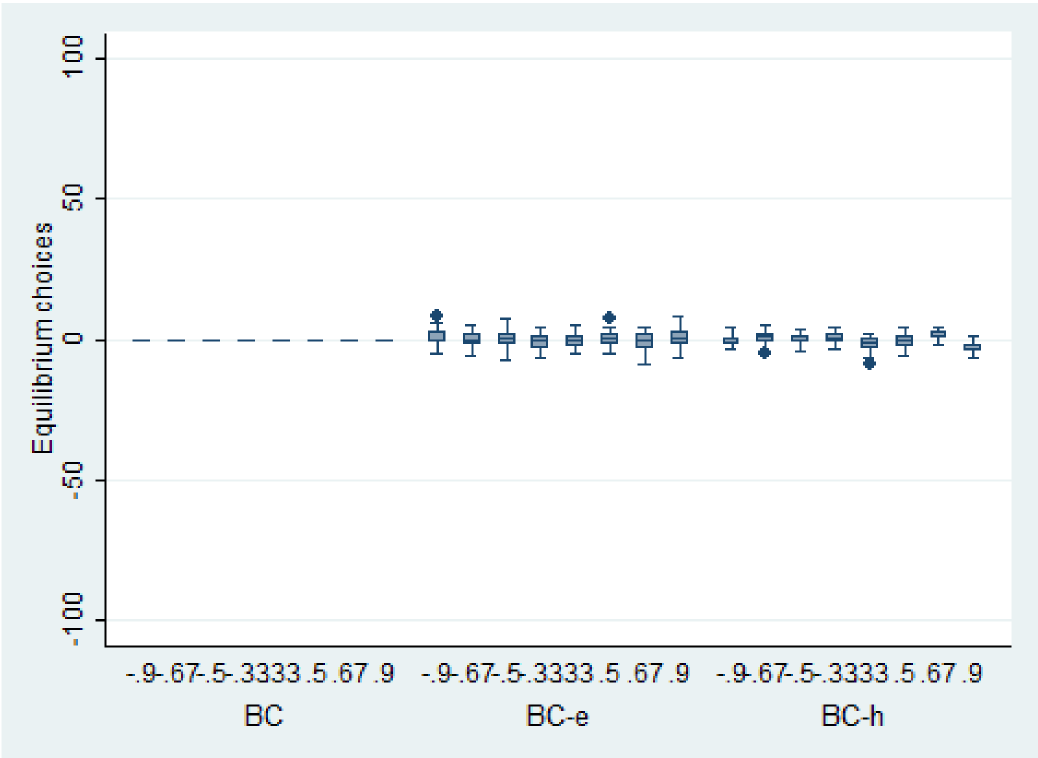

Figure A.1: Box-plots-Equilibrium choices for each possible game BC, BC-e, and BC-h, separately for each β-parameter.

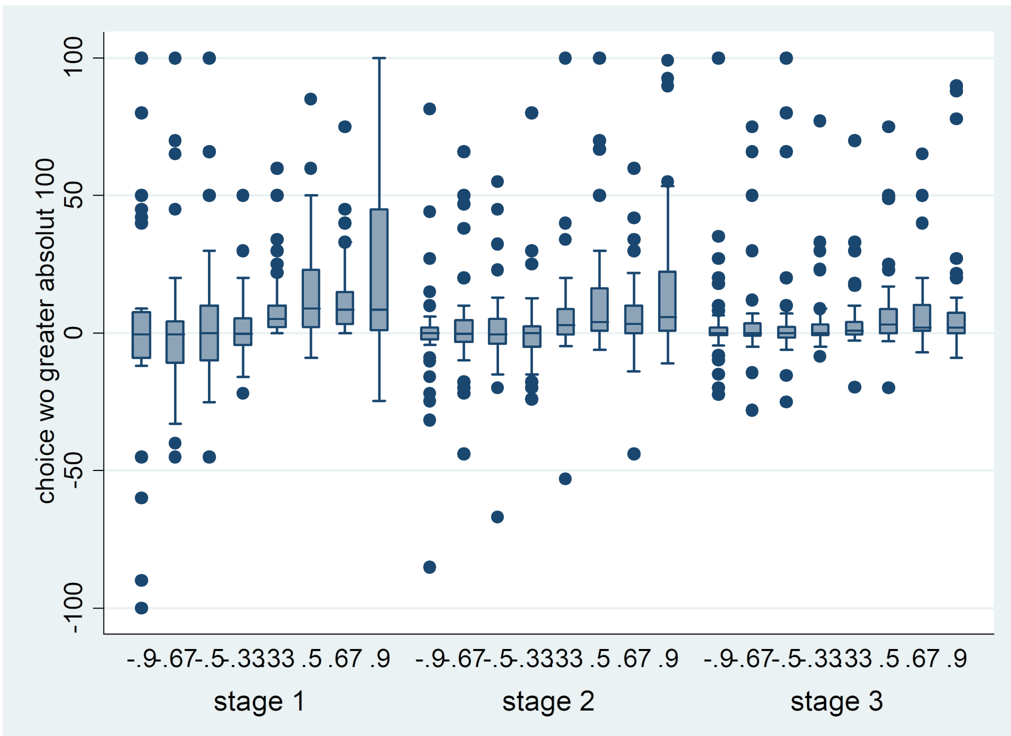

Figure A.2: Box-plots choices Order 1: stage 1 with BC, stage 2 with BC-e, stage 3 with BC-h, separately for each β-parameter within each stage. Choices greater than 100 in absolute value are excluded.

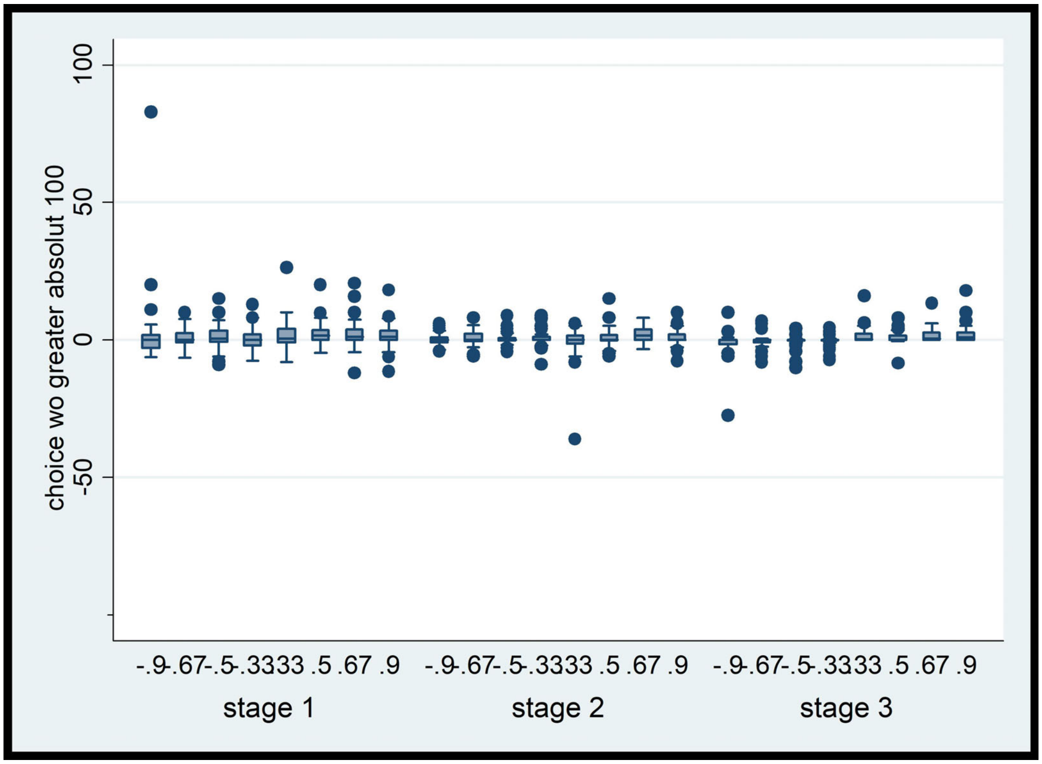

Figure A.3: Box-plots choices Order 2: stage 1 BC-e, stage 2 BC-h, stage 3 BC, separately for each β-parameter within each stage. Choices greater than 100 in absolute value are excluded.

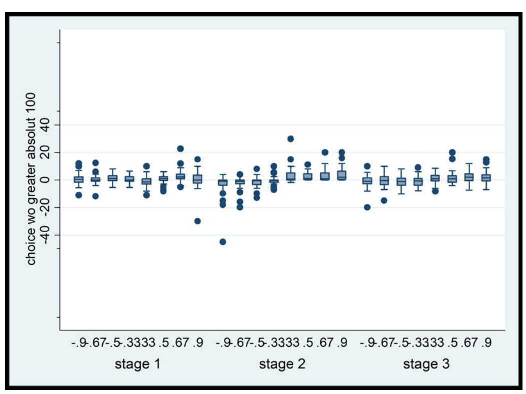

Figure A.4: Boxplots choices Order 3: stage 1 BC-h, stage 2 BC-, stage 3 BC-e, separately for each β-parameter within each stage. Choices greater than 100 in absolute value are excluded.

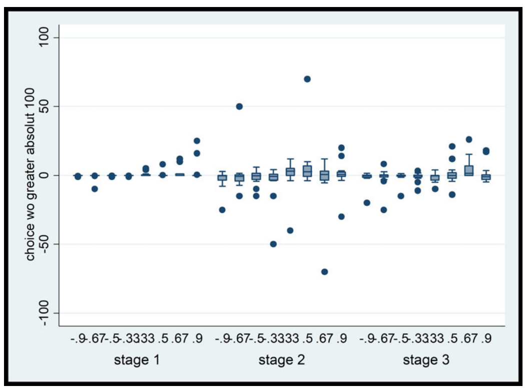

Figure A.5: Box-plots choices of 18 Professionals with order 1: stage 1 BC, stage 2 BC-e, stage 3 BC-h, separately for each β-parameter within each stage. Choices greater than 100 in absolute value are excluded.

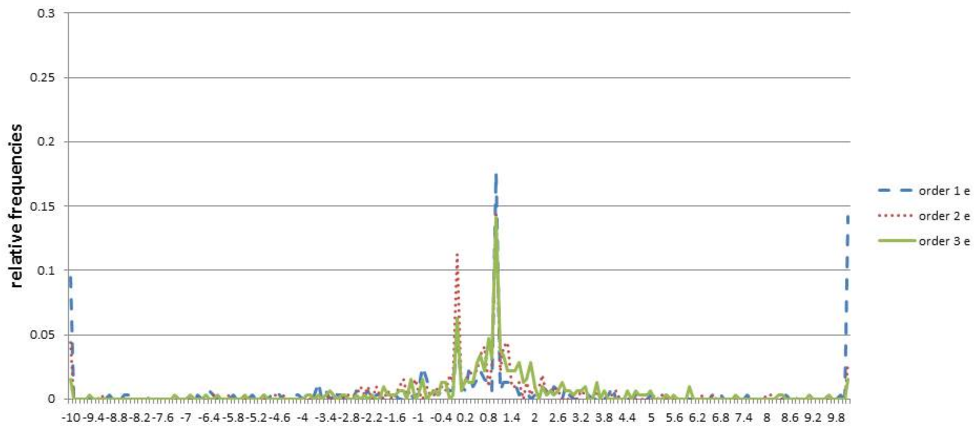

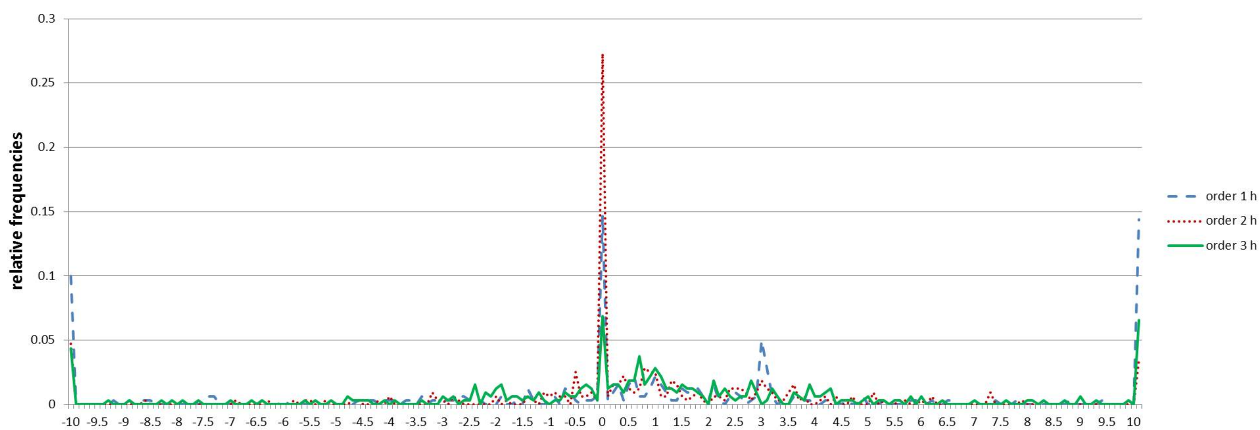

Figure A.6: Upper figure: Relative frequency of γ coefficients for all BC-e games (pooled over different β values), shown separately for each treatment order. There are "fat tails" especially when BC-e is played after BC (in order 1). Modal coefficients are at 1 (equilibrium signal). But there are also spikes at 0 (ignore signal, especially in order 2 and 3). Lower figure: similarly reports frequencies of γ coefficients for all BC-h games. There are modes at 0 (ignoring signals) or 3 (ignoring the common signal) and "fat tails" when BC-h is player after BC (in order 1). There is no spike at the BC-h equilibrium, $\gamma = 1.5$ Each line in these Figures contains $8 \times 40 = 320$ observations.

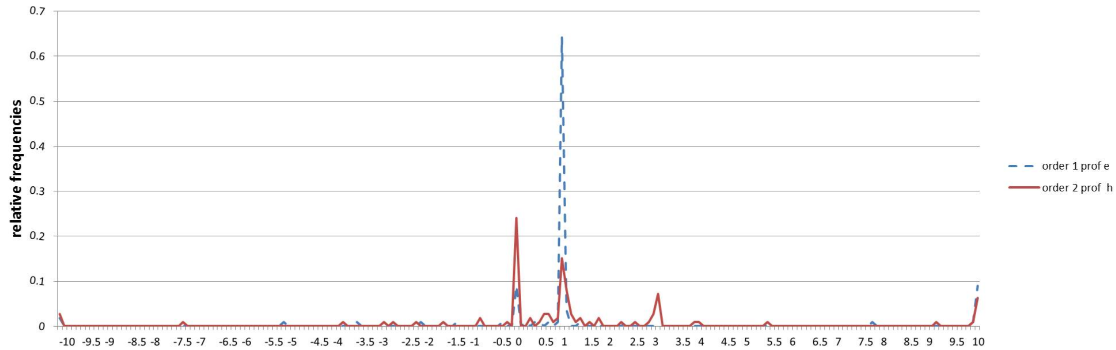

Figure A.7: Relative frequency of γ coefficient for 18 Professionals separately for BC-e and BC-h. Only order 1 is played. We see again that in the BC-h game with professionals, there is no spike at the BC-h equilibrium, $\gamma = 1.5$.

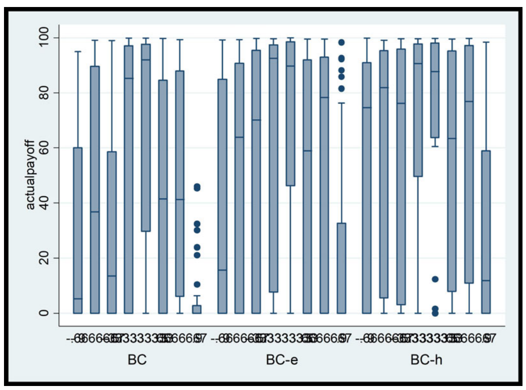

Figure A.8: Payoffs (=efficiency) in order 1. One can see a wide payoff range because of the large number of outliers.

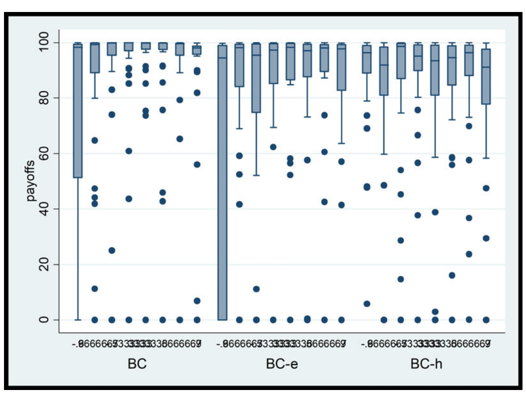

Figure A.9: Payoffs (efficiency) in order 2; note that the games are NOT sorted by stages but the games BC (stage 3), BC-e (stage 1), BC-h (stage 3) for easier comparison between figure A 8-10.

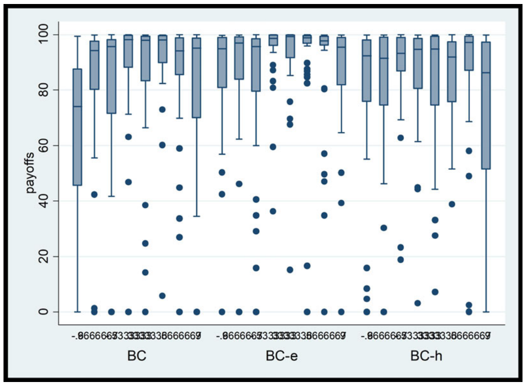

Figure A.10: Payoffs (efficiency) in order 3; note that the games are NOT sorted by stages but the games BC (stage 2), BC-e(stage 3), BC-h (stage 1) for easier comparison between figure A 8-10.

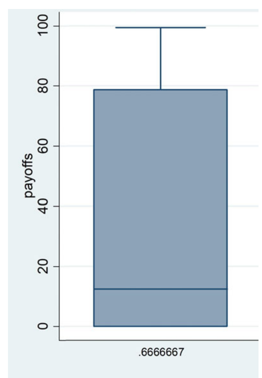

Figure A.11: Payoffs (efficiency) in the 2/3 BC game with bounded interval 0 to 100;

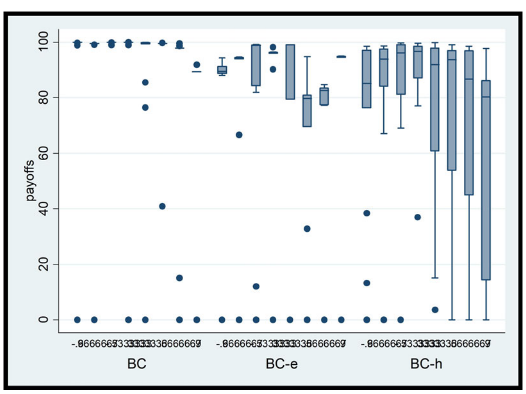

Figure A.12: Professionals playing order 1. Note that in the BC-h treatment payoffs also show a large variance.