Weizsäcker, Georg

Ignoring the rationality of others: evidence from experimental normal-form games

Article (Accepted version)
(Refereed)

Original citation:
Weizsäcker, Georg (2003) Ignoring the rationality of others: evidence from experimental normal-form games. *Games and economic behavior*, 44 (1). 145 -171.
DOI: 10.1016/S0899-8256(03)00017-4

Ignoring the rationality of others: Evidence from experimental normal-form games[¤]

by

Georg Weizsäcker

Department of Economics, Harvard University, and Harvard Business School[y]

Suggested running title: Ignoring the rationality of others

[¤]I wish to thank Jörg Breitung, Miguel Costa-Gomes, Werner Güth, Ste¤en Huck, Dorothea Kübler, Jörg Oechssler, Al Roth, two anonymous referees, and several audiences for their helpful comments, and Miguel Costa-Gomes and Dale Stahl for providing data and experimental instructions.

[y]Mailing address: Harvard Business School, Sherman Hall, Boston, MA 02163, USA. Email address: weizsack@fas.harvard.edu. Phone: +1-617-4951810. Fax: +1-617-4962203.

Abstract

Two behavioral models of two-person normal-form game play are presented and estimated, using three experimental data sets. The models are variants of the Quantal Response Equilibrium model de…ned by McKelvey and Palfrey (1995, *Games and Economic Behavior*), but allow a player to hold inaccurate beliefs about the behavior of her opponent. Each model involves two parameters: One captures the player's own level of response rationality, the other the level she attributes to her opponent. In order to allow for type heterogeneity among the subjects in the experiments, parametric distributions of these parameters are assumed. The estimation results indicate that the subjects' choices follow a speci…c anomalous pattern: On average, subjects play as if they signi…cantly underestimated their opponent's rationality.

JEL codes: C23, C91. Keywords: beliefs, prediction accuracy, experiments

1 Introduction

In analyses of game play data, the question often arises whether agents can be assumed to hold beliefs about their opponents' behavior that are on average correct. Particularly if one allows for boundedly rational decisions (or random perturbations in the subjects' utilities), the assumption that all players are perfectly informed about their opponents' propensities to err may not necessarily be satisfied. The interpretation of observed decisions, however, may crucially rely on whether or not this assumption is made.

In this paper, two closely related models of two-person normal-form game play are formulated and estimated, both of them allowing for relatively general sets of beliefs about the opponent's behavior. The models are based on the Quantal Response Equilibrium (QRE) defined by McKelvey and Palfrey (1995), who proposed a statistical reaction function prescribing the players' choice probabilities. In contrast to the standard QRE analysis, the behavioral models used in this paper each involve a second response function which a player attributes to her opponent's play, without restricting the actual choices of the opponent to follow this perceived behavior. For example, both players of a two-person game are allowed to view their opponents as highly irrational agents without responding irrationally to these beliefs themselves. Or, the models also allow the agents to perceive their opponents as behaving more rationally than they do themselves (where rationality is understood as precision of best responding to given beliefs). It can therefore be tested whether subjects in laboratory experiments systematically mispredict the degree of their opponents' rationality and what kinds of mispredictions the subjects exhibit. For the data considered below, the estimation results suggest a consistent pattern of subject behavior: On average, subjects choose as if they significantly underestimate the rationality of their opponents (or simply tend to ignore the other player's choices) and relatively consistently play best responses against these beliefs.[1]

Several models of normal-form game play that also focus on the beliefs subjects hold about their opponents' actions have been estimated and tested in the experimental literature. Stahl and Wilson (1995) formulate a multiple-type model, following Nagel's (1995) step-j-thinking approach, in which each player is assumed to belong to one of several types of players: A 'level-0' type chooses randomly with equal probabilities over his actions, a 'level-1' type consistently plays a best response to level-0 behavior (i.e., views her opponents to choose at random), a 'level-2' type best responds to a mixture of level-0 and level-1 types, a Nash type chooses Nash Equilibrium strategies, a 'worldly' type responds to a hypothetical mixture of level-0, level-1, and Nash types, and a 'rational expectations' type correctly anticipates the proportions of all types in the subject pool and chooses a best response to the resulting

[1] To be more precise, one of the two models produces this result for all three data sets. Under the assumptions of the other model, the analogous result shows up for two of the three data sets, with a reverse (and weakly significant) result for the third.

probability distribution over the opponent's action set. The model is estimated using a set of 3x3 normal-form games, resulting in estimated proportions of about 17% of the subjects being level-0 types, 20% level-1 types, 2% level-2 types, 17% Nash types, 43% worldly types, and 0% rational expectations types. In Stahl and Wilson (1994) and Costa-Gomes et al. (2001), related models are estimated from di¤erent sets of normal-form game data. Their results, however, suggest quite di¤erent type distributions in the subject pools.[2]

In the context of the present study, at least two conclusions can be drawn from these experimental papers. First, there appears to be a large variance in the subjects' behavioral patterns and in their beliefs about the opponents' play. No single type of player can account for the greater part of the observed behavior in the three experiments. (Rather, the estimation results seem to be highly sensitive toward the underlying sets of player types in the models and/or toward the speci…c games used in the experiments.) Second, both the Nash type and the level-1 type constitute non-negligible parts of the subject populations in all three experiments. In terms of beliefs attributing a certain degree of rationality to one's opponents, both extremes seem to

[2] Stahl and Wilson (1994), whose model includes only the level-0, level-1, level-2, and Nash types, report estimates of 24% level-1 types, 49% level-2 types, 27% Nash types, and an insigni…cant proportion being level-0 types. The model by Costa-Gomes et al. (2001) includes a number of additional types, but also the level-1 and Nash types, which are estimated to account for 23% and 8% of the subject pool, respectively.

be present in the subject pools.[3] Given this evidence, the question adressed here is whether the extent to which subjects implicitly rely on the opponent's rationality does, in the aggregate, re‡ect the actual response rationality (with which subjects respond to their beliefs). This question cannot be answered within the multiple-type models cited above, because in the construction of the player types' beliefs, some of the types are left out of consideration, so beliefs cannot coincide with actual behavior.

In contrast, the behavioral models de…ned below make use of the symmetric structure of the Quantal Response Equilibrium, in the sense that a player is aware of the fact that her opponent follows the same behavioral model as herself. At the same time, since the models allow for inaccurate perceptions of the opponent's response precision, the "rational expectations" assumption made in previous QRE analyses is relaxed and can be tested. To this end, the two models each involve two parameters in their basic formulations: One parameter represents the response precision that a player has herself, the other the precision level that she attributes to her opponent.

[3] For further results supporting this observation see the games conducted by Van Huyck et al. (1990, 1991), Beard and Beil (1994), Schotter et al. (1994), Nagel (1995), Ho et al. (1998), Haruvy and Stahl (1998), and Stahl (1999). The question remains, however, whether the observed choices can in fact be seen as responses to explicit beliefs held by the subjects, or whether the apparent belief distortions are due to decision heuristics that do not in any consistent way involve responding to beliefs about one's opponents. On this issue see Huck and Weizsäcker (2002) who use direct methods of belief elicitation to con…rm the apparent level-1 behavior in experimental games.

The two models di¤er only with respect to the error structure that is assumed, one using a logistic error speci…cation, the other uniform errors. Both models include as special cases the level-0, level-1, and Nash types of Stahl and Wilson (1995), as well as a continuum of other types. In order to allow for variance in beliefs and in precision of responses to given beliefs, the models' parameters are assumed to be distributed according to parametric distributions (with given functional forms), and the distributions are estimated using maximum-likelihood techniques. The resulting two econometric models are comparably ‡exible in terms of allowing for a variety of beliefs and response behavior, and they each involve four parameters that are to be estimated from the data.

The experimental data sets used in the statistical analysis are taken from the studies by Stahl and Wilson (1994, 1995) and Costa-Gomes et al. (2001).[4] The estimation results show that the average belief that subjects hold about their opponent's response precision is signi…cantly below the actual average response precision. This can be interpreted as a failure of the "average" subject to predict her opponent's rationality accurately: Subjects act as if they tend to ignore their opponents' rationality, and could earn more in the experiments if they did not do so.

The remainder of the paper is organized as follows: The next section formulates

[4] Some of the games are not appropriate for estimating the models, so not all of the data will be used (see Section 3).

the general behavioral model and the two special cases that will be used in the statistical analysis. An example and a short discussion accompany the definitions. Section 3 presents the maximum-likelihood estimation. In order to give an impression of the robustness of the results, the estimations are also conducted separately for the three data sets, and for a number of other subsets of the experimental games. Section 4 concludes.

2 Two models of normal-form game play

Consider a finite two-player game $\Gamma = \{a, b, A^a, A^b, u^a, u^b\}$ with players $i = a, b$, where A^i is player i's action set and $u^i(\cdot)$, $u^i : A^a \times A^b \to \Re$, is player i's (von Neumann-Morgenstern) utility function, which is assumed to be bounded. Let $m^i = \# A^i$ be the number of actions i has to choose from. Furthermore, let Q^i be the set of probability distributions over A^i, i.e., Q^i is player i's set of mixed strategies, where, for any $\sigma^i \in Q^i$, $\sigma^i(a^i)$ denotes the probability that $a^i \in A^i$ is chosen according to σ^i.

The basic concept of the subsequent analysis is to make use of response functions attributed to the players. Define the actual response function of player i, denoted by $r^i(\cdot)$, $r^i : Q^j \to Q^i$, as the mixed strategy that i chooses in response to a mixed strategy $\sigma^j \in Q^j$ of player j. For example, $r^i(\sigma^j)$ could be an element of the best response correspondence of i against σ^j, defined by

$$BR^i(\sigma^j) = \arg\max_{\sigma^i} u^i(\sigma^i, \sigma^j), \quad i, j \in \{a, b\}, i \neq j.$$

Or, $r^i(\pi^j)$ could be a constant function not depending on π^j, e.g. the uniform distribution over all m^i actions, which will be denoted by

$$P_0^i = (\frac{1}{m^i}, ..., \frac{1}{m^i}), \quad i \in \{a, b\}.$$

To allow for general beliefs about the opponent's strategy, define also the perceived response function of player j, $\rho^j(\cdot)$, $\rho^j : Q^i \to Q^j$, as the belief that i holds about the response function employed by j, with $i \neq j$. In analogy to the above examples, $\rho^j(\pi^i)$ could e.g. be an element of $BR^j(\pi^i)$ or be equal to the constant function P_0^j.

Using these two response functions, a slight variation of McKelvey and Palfrey's definition of a QRE strategy can be given as follows.[5]

Definition 1 For a given pair of response functions $(r^i(\cdot), \rho^j(\cdot))$, player i's strategy π^i is a Quantal Response Equilibrium strategy (QRE strategy) if

$$\pi^i = r^i(\rho^j(\pi^i)).$$

If player i plays a QRE strategy for $(r^i(\cdot), \rho^j(\cdot))$, she assumes that j acts according to $\rho^j(\cdot)$ and responds herself to this belief according to $r^i(\cdot)$, such that the above

[5] Definition 1 and the definition of a QRE strategy in McKelvey and Palfrey (1995) are not equivalent because in their formulation the response functions have to be derived from maximization in a random-utility environment. This theoretically more satisfying approach is, however, much more "definitionally intensive". The special cases of the QRE strategies used in the data analysis below will be accompanied by random-utility justifications, one of them strictly along the lines of McKelvey and Palfrey (1995).

fixed-point property holds. For example, if $(r^i(\cdot), \tilde{r}^j(\cdot))$ are best responses of the two players, then a QRE strategy π^i is a Nash Equilibrium strategy. Notice that the actual behavior of player j does not enter the definition of a QRE strategy – rather it is only i's perception of j's behavior that matters. The approach is entirely decision-theoretic, because the QRE is understood here as an equilibrium as perceived by player i: The fixed-point property above makes a choice prediction for only one player, i, who behaves as if both players play according to a QRE with response functions $(r^i(\cdot), \tilde{r}^j(\cdot))$. In the data analysis, the fact that QRE strategies are defined with respect to one player only (in the version of Definition 1) will allow the case that the predictions for the two players of a game are inconsistent.

In order to obtain well-defined choice predictions that can be analyzed using data from experimental games, it remains to specify the form of the response functions $r^i(\cdot)$ and $\tilde{r}^j(\cdot)$. Before turning to the data analysis in Section 3, the following subsections therefore present two parametrized special cases of QRE strategies.

2.1 Asymmetric Logit Equilibrium strategies

First define $\bar{u}^i_k(\pi^j)$ to be player i's expected utility from playing the (pure) action a^i_k against the strategy π^j of player j; i.e., $\bar{u}^i_k(\pi^j) = \sum_{l'=1}^{m^j} \pi^j(a^j_{l'})\, u^i(a^i_k, a^j_{l'})$; $i, j = a, b$. Player i's actual response function $r^i(\cdot)$ is the logistic response function with precision parameter λ^i ($\lambda^i \geq 0$) if for all actions $a^i_k \in A^i$ and all player-j strategies $\pi^j \in Q^j$ it holds that the probability weight on the action a^i_k, denoted by $r^i_k(\pi^j;\lambda^i)$, is given by

$$r^i_k(\pi^j;\lambda^i) = \frac{e^{\lambda^i \bar{u}^i_k(\pi^j)}}{\sum_{k'=1}^{m^i} e^{\lambda^i \bar{u}^i_{k'}(\pi^j)}}.$$

The logistic response function has the property that the probability weights on the actions a^i_k are ordered in correspondence to the ordering of the expected utilities $\bar{u}^i_k(\pi^j)$, $k = 1;\ldots;m^i$. Also, varying the precision parameter λ^i corresponds to varying the distance between the logistic response and the best response to π^j. The greater λ^i, the more probability weight lies on the actions that are best responses to π^j. As λ^i approaches infinity, $r^i_k(\pi^j;\lambda^i)$ approaches zero if and only if a^i_k is not a best response to π^j.

Analogously, define $\bar{u}^j_l(\pi^i)$ to be j's expected utility from action a^j_l chosen in response to strategy π^i. Then, the perceived response function $\tilde{r}^j(\cdot)$ is the logistic response function with precision parameter $\tilde{\lambda}^j$ ($\tilde{\lambda}^j \geq 0$) if for all $a^j_l \in A^j$ and all $\pi^i \in Q^i$ it holds that the probability weight on a^j_l is given by

$$\tilde{r}^j_l(\pi^i;\tilde{\lambda}^j) = \frac{e^{\tilde{\lambda}^j \bar{u}^j_l(\pi^i)}}{\sum_{l'=1}^{m^j} e^{\tilde{\lambda}^j \bar{u}^j_{l'}(\pi^i)}}.$$

Definition 2 For given precision parameters λ^i and $\tilde{\lambda}^j$, let $r^i(\cdot)$ and $\tilde{r}^j(\cdot)$ be the logistic response functions with λ^i and $\tilde{\lambda}^j$ respectively. Then player i's strategy π^i is an Asymmetric Logit Equilibrium strategy (ALE strategy) with parameters $(\lambda^i;\tilde{\lambda}^j)$ if π^i is a QRE strategy for $(r^i(\cdot);\tilde{r}^j(\cdot))$.

The class of ALE strategies allows for a relatively wide variety of game play behavior. For example, in the special case that λ^i and $\tilde{\lambda}^j$ are equal, $\lambda^i = \tilde{\lambda}^j = \lambda$, the set of ALE strategies with parameters $(\lambda^i; \tilde{\lambda}^j)$ is equal to the set of Logit Equilibrium strategies with parameter λ, as defined by McKelvey and Palfrey (1995). Hence, any Logit Equilibrium strategy is an ALE strategy. As with Logit Equilibrium strategies, a player takes the opponent's propensity to choose non-optimal responses into consideration when playing an ALE strategy (as well as the fact that the opponent does so, too, and so on) and chooses in accordance with the solution of the resulting fixed-point problem. However, the difference between the two choice predictions is that playing an ALE strategy a player with a precision parameter λ^i is not assumed to attribute the identical parameter value to her opponent's behavior. The parameter $\tilde{\lambda}^j$ only reflects i's expectation of j's behavior, and no consistency or "rationality" of expectations will be imposed in the data analysis below.[6]

Also, the class of ALE strategies encompasses most of the player types introduced

[6] Player i's second-order belief concerning the response precision, i.e. the belief she supposes her opponent to hold about λ^i, is assumed to be correct, as in the Logit Equilibrium. Although this assumption could, in principal, be relaxed as well, such a more general analysis is not attempted here. One may argue that it is more natural to assume that second-order beliefs coincide with the "true" precision level λ^i than to assume that first-order beliefs ($\tilde{\lambda}^j$) do, because second-order beliefs concern the player's own behavior. However, it is clearly possible that a misspecification of second-order beliefs leads to the apparent distortions of first-order beliefs that are reported in Section 3. For alternative specifications of second-order (and higher-order) beliefs in models of quantal response see Stahl and Wilson (1994, 1995), Goeree and Holt (2000) and Kübler and Weizsäcker (2002).

by Stahl and Wilson (1994, 1995). In particular, if $\lambda^i = 0$ holds, then i is a level-0 type. If $\lambda^i = 1$ and $\tilde{\lambda}^j = 0$, i is a level-1 type. In the case $\lambda^i = \tilde{\lambda}^j = 1$, i is a Nash type. Hence, the class of ALE strategies can be seen as a two-dimensional continuum between these three archetypical patterns of behavior. For intermediate values of λ^i and $\tilde{\lambda}^j$, player i exhibits a 'worldly-like' behavior: She considers j to be neither completely irrational nor unboundedly precise in his responses, and she takes into account the fact that j believes her, i, not to be perfectly rational either.[7]

As an illustration, consider the 2x2 game Γ_1 given in the Table I, which is one of the games used for the data analysis in Section 3. In Γ_1, the column player's strategy R is dominated, and the game has a unique Nash Equilibrium at $(U; L)$.

Insert Table I about here.

The ALE strategies for different parameter constellations $(\lambda^{row}; \tilde{\lambda}^{col})$ are shown in Figure 1.[8] Figure 1a depicts the graph of the row player's ALE strategies for

[7] Due to the symmetric structure of the ALE strategy, player i is assumed to expect her opponent j to exhibit this 'worldly-like' behavior, too, if $(\lambda^i; \tilde{\lambda}^j) > 0$. In contrast, the worldly type of Stahl and Wilson (1995) expects level-0, level-1, and Nash behavior but does not believe in the existence of other worldly players.

[8] The strategies were calculated with a grid-search algorithm similar to the "Gobit-All" algorithm used in the software package Gambit, which produces QRE strategies (see McKelvey et al., 1996). The entries in the matrix were transformed into US$ payoffs in the same way as in the experiment by Costa-Gomes et al. (2001), where the game was used; see Appendix B.

varying values of $\tilde{\lambda}^{col}$: λ^{row} is held constant at $\lambda^{row} = 16$, and $\tilde{\lambda}^{col}$ varies on a logarithmic scale. In Figure 1b, $\tilde{\lambda}^{col}$ is held constant at $\tilde{\lambda}^{col} = 16$, and λ^{row} varies. Both parts of the figure show the ALE probability of the row player choosing action U (marked "Pr(U)") and the row player's corresponding belief about the column player's probability to choose L (marked "Pr(L)").

Insert Figure 1 about here.

Figure 1a illustrates the dependence of the row player's behavior on the precision parameter $\tilde{\lambda}^{col}$ that she attributes to the column player. For small values of $\tilde{\lambda}^{col}$, implying in a perceived strategy of the column player that is close to the uniform distribution, the row player "rationally" responds with almost none of the probability weight lying on U. As $\tilde{\lambda}^{col}$ increases, i.e., as the column player's hypothetical response becomes closer to the best response, the row player increases her probability weight on U, and eventually (for large values of $\tilde{\lambda}^{col}$) her strategy approaches the Nash Equilibrium strategy. Figure 1b, on the other hand, shows that regardless of the value of her own precision parameter λ^{row}, the row player expects her opponent to choose L (which is his best response) with almost full probability mass, so she responds by playing U with the greater probability. Again, for large values of λ^{row}, i.e., as her own response precision gets larger, her ALE strategy approaches the Nash Equilibrium strategy.

Viewing the game from the perspective of the column player, notice that due to the symmetric structure of the QRE condition (see Definition 1) it holds that whenever π^i is an ALE strategy for player i with parameters $(\lambda^i; \tilde{\lambda}^j)$ it is also true that $\pi^j = \tilde{\rho}^j(\pi^i; \tilde{\lambda}^j)$ is an ALE strategy for j with parameters $(\tilde{\lambda}^j; \lambda^i)$. Hence, the column player's hypothetical choice probabilities depicted in the figure can equally be viewed as his actual choice probabilities prescribed by ALE strategies with the parameter constellations in the reverse order.

Following this interpretation, Figure 1b also shows that the column player's actual ALE behavior with a precision level of $\lambda^{col} = 16$ does not depend on his belief about the row player to any significant degree. Since R is a dominated strategy, it is always a best response for the column player to choose L. Similarly, in Figure 1a, as λ^{col} rises (for a given value of $\tilde{\lambda}^{row}$), more and more probability mass lies on L, and the column player's ALE strategy approaches the Nash Equilibrium strategy.[9]

[9] An ALE strategy can alternatively be viewed as the equilibrium choice distribution in a game with random errors (following extreme-value distributions) perturbing the payoffs. For Logit Equilibrium strategies, this was demonstrated by McKelvey and Palfrey (1995). Following their approach, one can immediately establish two important properties of ALE strategies: First, ALE strategies exist for all values of $(\lambda^i; \tilde{\lambda}^j)$, and second, if λ^i and $\tilde{\lambda}^j$ both converge to infinity, any limiting ALE strategy is a Nash Equilibrium strategy of the unperturbed game. These results follow from McKelvey and Palfrey (1995), Theorem 1 and Theorem 2, respectively. For a discussion see also McKelvey and Palfrey (1998) as well as Zauner (1999), who applies different specifications of the error structure to random-utility models. A common interpretation of such random-utility perturbations is to view

2.2 Asymmetric Noisy Nash Equilibrium strategies

This subsection presents a parametrization of QRE strategies that is, although somewhat less elegant, easier to compute and to interpret than the one in the previous subsection. This second model is the uniform-error analogon to the class of ALE strategies, meaning that the players are assumed to erroneously choose non-optimal responses with constant probabilities, which do not depend on the expected utilities resulting from these actions. Specifically, some proportion of the probability weight in a player i's response function is assumed to lie on the uniform distribution over all of her m^i actions, P_0^i.

Player i's actual response function $r^i(\cdot)$ is called the uniform-error response function with error rate ε^i ($\varepsilon^i \in [0;1]$) if, for all player-j strategies $\pi^j \in Q^j$,

$$r^i(\pi^j;\varepsilon^i) = (1-\varepsilon^i)\widehat{BR}^i(\pi^j) + \varepsilon^i P_0^i;$$

where $\widehat{BR}^i(\pi^j)$ is an arbitrary element of i's set of best responses to π^j, $BR^i(\pi^j)$, and $i;j \in \{a;b\}; i \neq j$. Analogously, player j's perceived response function $\hat{r}^j(\cdot)$ is called the uniform-error response function with error rate $\hat{\varepsilon}^j$ ($\hat{\varepsilon}^j \in [0;1]$) if for all $\pi^i \in Q^i$ it holds that

$$\hat{r}^j(\pi^i;\hat{\varepsilon}^j) = (1-\hat{\varepsilon}^j)\widehat{BR}^j(\pi^i) + \hat{\varepsilon}^j P_0^j;$$

where $\widehat{BR}^j(\pi^i) \in BR^j(\pi^i)$, $i;j \in \{a;b\}; i \neq j$.

---

them as representing the impact of computational errors.

Definition 3 For a given pair of error rates $(\epsilon^i;\hat{\epsilon}^j)$, let $r^i(\cdot)$ and $\hat{r}^j(\cdot)$ be the uniform-error response functions with ϵ^i and $\hat{\epsilon}^j$ respectively. Then player i's strategy π^i is an Asymmetric Noisy Nash Equilibrium strategy (ANNE strategy) with parameters $(\epsilon^i;\hat{\epsilon}^j)$ if π^i is a QRE strategy for $(r^i(\cdot);\hat{r}^j(\cdot))$.

As with ALE strategies, the class of ANNE strategies encompasses a variety of behavioral patterns, depending on the pair of parameters $(\epsilon^i;\hat{\epsilon}^j)$. If $\epsilon^i = 1$, then i is a level-0 type and fully randomizes between her actions. If $\epsilon^i = 0$ and $\hat{\epsilon}^j = 1$ hold, i is a level-1 type, and in the case $\epsilon^i = \hat{\epsilon}^j = 0$ she is a Nash type. For intermediate values of $(\epsilon^i;\hat{\epsilon}^j)$, player i plays 'worldly-like', similar to the case of an ALE strategy with $(\lambda^i;\hat{\lambda}^j) > 0$: She thinks of her opponent as neither perfectly precise nor perfectly imprecise in his responses, and she considers the fact that j also takes her own imprecision into account.

In contrast to ALE strategies, which cannot be derived analytically for most games and parameter pairs $(\lambda^i;\hat{\lambda}^j)$, ANNE strategies can be calculated exactly, in analogy to solving for the Nash Equilibrium strategies of a game. Figure 2 depicts the ANNE strategies of game Γ_1 (see Table I) for varying parameter constellations, and for both players (again marked "Pr(U)" and "Pr(L)", indicating the respective probabilities to choose the Nash strategies). In Figure 2a, ϵ^{row} is held constant at $\epsilon^{row} = 0.2$ and $\hat{\epsilon}^{col}$ varies between 0 and 1, whereas in Figure 2b ϵ^{row} varies and $\hat{\epsilon}^{col}$ is held constant at 0.2.

Insert Figure 2 about here.

As Figure 2a shows, the row player's ANNE strategy is close to her Nash Equilibrium strategy U as long as $\hat{\lambda}^{col}$ is below some critical value $\hat{\lambda}^{col}_{crit} \approx 0.73$. For values of $\hat{\lambda}^{col}$ above $\hat{\lambda}^{col}_{crit}$, the row player considers the column player to randomize enough to make the action D her best response against any resulting probability distribution. As is the case with ALE strategies, the column player's perceived choice probabilities depicted in the figures can equally be viewed as his actual ANNE strategies with the reversed order of parameters. Hence, Figure 2b shows that under the ANNE assumptions, paralleling the ALE case, the column player plays his dominant strategy L with a high probability regardless of the row player's actions.

A comparison of Figures 1a and 2a shows that the change in the row player's behavior is smoother with ALE strategies and varying λ^{col} than it is with ANNE strategies and varying $\hat{\lambda}^{col}$. Hence, the ALE strategy model is more flexible in the sense that even for a high precision in her actual responses, the row player may choose "intermediate" strategies.[10]

[10] To explore the applicability of the ANNE strategy model, it is useful to ask for a random-utility justification of the assumption that players choose ANNE strategies, similar to the reinterpretation of ALE strategies in Footnote 9. Using standard concepts of Bayesian game analysis, it is possible, for any pair of parameters $(\varepsilon^i, \hat{\lambda}^j) \in [0,1] \times [0,1]$, to construct a finite game of incomplete information in which the players' equilibrium probability distributions over their choices are equal to the probability distributions resulting from playing ANNE strategies with parameters $(\varepsilon^i, \hat{\lambda}^j)$. This construction is

3 Data analysis

3.1 Likelihood functions

In this subsection, the basic assumptions underlying the parameter estimations are formulated – in particular concerning the introduction of type heterogeneity – and likelihood functions for both the ALE strategy model and the ANNE strategy model are presented. The following subsections will describe the experimental data, and present and illustrate the estimation results.

Consider a number of N subjects, $i = 1; :::; N$, each of whom is confronted with a set of H^i two-person normal-form games. Let $c(i; h)$ be the action chosen by subject i in game h, $h = 1; :::; H^i$. Also, let $c^i = fc(i; h)g_{h=1}^{H^i}$ denote the vector of i's joint choices.

Now suppose that choice probabilities are given by a model of normal-form game play, such as the ALE strategy model or the ANNE strategy model. Let $P_{ihk}(\mu^i)$ be the probability that subject i chooses action k in game h, where μ^i is i's vector

---

relegated to Appendix A. Due to this reinterpretation, existence of ANNE strategies is established. Also, notice that it holds for almost all games ¡ and all parameters $(^2{}^i; \hat{e}^j)$ that the ANNE strategy for a player is unique if either she or her opponent has a strictly dominant strategy, because the other player then has a (generically) unique best response. This property is useful for the data analysis in the following section, in that it provides a simple sufficient condition for the model to make an unambiguous prediction.

of parameters determining her choice probabilities in the underlying model: In the context of ALE strategies, θ^i is given by $(\lambda^i; \hat{\lambda}^j)$; in the context of ANNE strategies, θ^i is given by $(\varepsilon^i; \hat{\varepsilon}^j)$. The vector θ^i can also be thought of as subject i's type. With this notation, the probability that i chooses $c(i;h)$ in game h is equal to $P_{ihc(i;h)}(\theta^i)$. Hence, the likelihood of subject i's type θ^i, given her joint choices c^i, can be written as

$$\hat{L}^i(\theta^i|c^i) = \prod_{h=1}^{H^i} P_{ihc(i;h)}(\theta^i). \tag{1}$$

An assumption that is implicitly made by the construction of $\hat{L}^i(\theta^i|c^i)$ is that subject i's type is fixed over all games h, $h = 1;\ldots;H^i$. (Given this assumption, it is an important feature of both ALE strategies and ANNE strategies that they allow a player i with a fixed type θ^i to exhibit a flexible behavior over different games, depending on the payoffs of the games.[11]) Also, it is assumed that the H^i choices are made independently, given the probabilities $P_{ihc(i;h)}(\theta^i)$.

In order to allow for heterogeneity among the N subjects, suppose that θ^i is a random variable drawn from a distribution with density function $g(\theta^i|\beta)$; where β is a vector of parameters determining the density function $g(\cdot)$. (The vector β will be

[11] Consider for example Figure 2a and imagine a second game, Γ_1', in which the critical value of $\hat{\varepsilon}^{col}$ at which the row player's strategy changes discontinuously, i.e. a value $\hat{\varepsilon}^{col'}_{crit}$ corresponding to $\hat{\varepsilon}^{col}_{crit}$, is lower: $\hat{\varepsilon}^{col'}_{crit} < \hat{\varepsilon}^{col}_{crit}$. Then for a continuum of parameter pairs $(\varepsilon^{row}; \hat{\varepsilon}^{col})$ the ANNE strategies of the row player would prescribe a high probability of choosing the Nash Equilibrium strategy in Γ_1, but not in Γ_1'. A similar observation applies to ALE strategies.

estimated from the data.) Subject i's contribution to the overall likelihood is then given by the expected value of $\hat{L}^i(\theta^i|c^i)$;

$$\begin{aligned} L^i(\beta|c^i) &= \int \hat{L}^i(\theta^i|c^i)g(\theta^i|\beta)d\theta^i \\ &= \int \prod_{h=1}^{H^i} P_{ihc(i,h)}(\theta^i)g(\theta^i|\beta)d\theta^i. \end{aligned}$$

Finally, let $c = \{c^i\}_{i=1}^N$ be the vector containing all observed choices by the N subjects. Assuming that the N subject types $(\theta^1, \ldots, \theta^N)$ are identically and independently distributed, with densities $(g(\theta^1|\beta), \ldots, g(\theta^N|\beta))$, the likelihood of the parameter vector β is

$$L(\beta|c) = \prod_{i=1}^{N} L^i(\beta|c^i). \tag{2}$$

It remains to specify the distribution of subject types, $g(\cdot)$. For the sake of greater parsimony of the models, standard parametric distribution functions, with given functional forms, are assumed.[12] In particular, suppose for the two parameters of the ALE strategies, λ^i and $\tilde{\lambda}^j$, that they are independently drawn from two gamma distributions with densities $f_\gamma(\lambda^i|\rho, K)$ and $f_\gamma(\tilde{\lambda}^j|\tilde{\rho}, \tilde{K})$, respectively, for all subjects $i = 1, \ldots, N$. A gamma distribution has a unimodal density function $f_\gamma(\cdot|\rho, K)$ de-

[12] An alternative approach would be to use nonparametric density estimations. Stahl and Wilson (1994, 1995), who base their analysis on a behavioral model that is related to ALE strategies and involves three type parameters, assume a distribution over the three-dimensional parameter space that contains several distinct mass points, reflecting the different player types.

pending on two distribution parameters, ρ and κ, and is defined by

$$f_\gamma(\lambda|\rho,\kappa) = \frac{\rho^\kappa}{\Gamma(\kappa)} e^{-\rho\lambda} \lambda^{\kappa-1}; \quad \lambda \geq 0; \rho > 0; \kappa > 0;$$

where $\Gamma(\kappa)$ denotes the gamma function given by

$$\Gamma(\kappa) = \int_0^\infty t^{\kappa-1} e^{-t} dt.$$

The expected values of λ^i and $\tilde{\lambda}^j$ are then given by $\frac{\kappa}{\rho}$ and $\frac{\tilde{\kappa}}{\tilde{\rho}}$ respectively, their variances are $\frac{\kappa}{\rho^2}$ and $\frac{\tilde{\kappa}}{\tilde{\rho}^2}$ respectively. Substituting the joint density of λ^i and $\tilde{\lambda}^j$,

$$g_\gamma(\lambda^i, \tilde{\lambda}^j|\rho,\kappa,\tilde{\rho},\tilde{\kappa}) = f_\gamma(\lambda^i|\rho,\kappa) \cdot f_\gamma(\tilde{\lambda}^j|\tilde{\rho},\tilde{\kappa}),$$

into (2) completes the construction of the likelihood function for the ALE strategy model.

The family of gamma distributions contains a number of familiar distributions as special cases, such as the χ^2-distribution, the exponential distribution, and the extreme value distribution. Also, it allows for the case that almost all probability mass is concentrated in an arbitrarily small interval around a specific value of λ^i (or $\tilde{\lambda}^j$ respectively), which corresponds to the case of homogeneity among the subjects. Hence, although the model contains only the four parameters $(\rho,\kappa,\tilde{\rho},\tilde{\kappa})$ that are to be estimated from the data, it allows for a relatively large variety of possible distributions of behavioral patterns in the experiment.

In the estimations using the ANNE strategy model, both parameters ε^i and $\tilde{\varepsilon}^j$ are assumed to be drawn from independent distributions that belong to the family of beta distributions over the range $[0, 1]$, with densities $f_\beta(\varepsilon^i|a, b)$ and $f_\beta(\hat{\varepsilon}^j|\hat{a}, \hat{b})$, respectively. A beta distribution depends on two distribution parameters a and b, and has a density function given by

$$f_\beta(\varepsilon|a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\,\Gamma(b)}\varepsilon^{a-1}(1-\varepsilon)^{b-1}, \quad \varepsilon \in [0, 1], \; a, b > 0,$$

where $\Gamma(\cdot)$ is defined as above. The beta density function is symmetric in the case $a = b$ and asymmetric otherwise, and it can be hump-shaped or U-shaped, depending on the values of a and b. Its mean is given by $\frac{a}{a+b}$, its variance is $\frac{ab}{(a+b+1)(a+b)^2}$. Like the gamma distribution, the beta distribution also contains the special case that almost all probability mass lies in an arbitrarily small interval around a specific value of ε.

For any subject i, the joint density of the actual and perceived error rates $(\varepsilon^i, \hat{\varepsilon}^j)$ is, under the above assumptions, given by

$$g_\beta(\varepsilon^i, \hat{\varepsilon}^j|a, b, \hat{a}, \hat{b}) = f_\beta(\varepsilon^i|a, b) \cdot f_\beta(\hat{\varepsilon}^j|\hat{a}, \hat{b}),$$

which can be substituted into (2).

3.2 Estimation results using all data

Stahl and Wilson (1994, 1995) conducted two experiments with 40 and 48 subjects, respectively, playing a total of 22 symmetric 3x3 normal-form games. Costa-Gomes, Crawford, and Broseta's (2001) experiment involved 72 subjects who were confronted with 18 normal-form (2x2, 3x2, and 4x2) games each. In the latter experiment, all 18 games were asymmetric and the subjects were divided into row players and column players, but the games were chosen such that the subjects were confronted with two almost identical versions of each of 9 games, played with switched player roles (and with slightly altered payo¤s), so all subjects played almost identical sets of games.[13] In all three experiments, no feedback information was given between the subjects' decisions. Also, importantly, in all sessions the experimenters took careful measures to ensure that the subjects understood the procedures and the rules of the games.[14]

Of the 40 games played, 23 have the property that both ALE strategies and ANNE strategies are unambiguous predictions for the greatest part of the possible parameter

[13] The data used are those of the "Baseline" and the "OB" treatment in the authors' terminology. One possible source of noise in the data is given by the fact that the 45 subjects of the Baseline treatment had to click on the cells of the payo¤ matrices with the computer mouse in order to see the payo¤s. However, Costa-Gomes et al. (2001) report statistical tests indicating that the choice behavior between the two treatments was di¤erent only within the limits of chance, so the data can be pooled. A further di¤erence in experimental design between the experiments occured concerning the payo¤ determination: While Costa-Gomes et al. (2001) randomly selected one opponent for each subject in each game, the payo¤s in the experiments in Stahl and Wilson (1994, 1995) were determined by calculating each participant's expected outcomes from playing against the distribution of all her possible opponents (i.e., subjects were "playing against the …eld") and by applying a binary-lottery procedure.

[14] Costa-Gomes et al. (2001) dismissed a total of 20 students who had shown up for the experiment, after they failed in the screening tests.

constellations $(\lambda^i; \tilde{\lambda}^j)$ and $(\epsilon^i; \tilde{\epsilon}^j)$ respectively.[15] More precisely, for these 23 games both the sets of ALE strategies and ANNE strategies are singletons for more than 95% of the possible parameter constellations (so a unique probability distribution over the players' actions is predicted for these parameters). In order to avoid equilibrium selection problems in the present analysis, only these 23 games are used for the estimations (Games 1, 2, 3, 5, and 8, in Stahl and Wilson, 1994, Games 1, 4, 5, and 12, in Stahl and Wilson, 1995, and all of Costa-Gomes et al.'s games except Games 5A, 6A, 7A, and 8A).[16] For all parameter constellations where more than one fixed point exists, it is, for simplicity, assumed in the data analysis that the subjects play each ALE strategy or ANNE strategy, respectively, with equal probability.

The 23 games used are shown in Appendix B, which also contains the aggregate choices of the $N = 160$ subjects. All of the selected games have unique Nash Equi-

---

[15] The parameter range for λ^i and $\tilde{\lambda}^j$ is artificially restricted to $[0; 30]$ in this and all subsequent parts of the analysis. For larger values of λ^i $[\tilde{\lambda}^j]$, the respective ALE strategies in all games are very close to the case that $\lambda^i = 30$ $[\tilde{\lambda}^j = 30]$. In particular, the ALE strategies, or Logit Equilibrium strategies, for $\lambda^i = \tilde{\lambda}^j = 30$ prescribe choice probabilities that are essentially identical to the Nash Equilibrium predictions, for the greatest part of the games.

[16] The relative numbers of unique predictions made by the models are calculated by solving for the ALE and ANNE predictions on a finite grid over the two-dimensional parameter spaces. The cutoff value of 95% is chosen somewhat arbitrarily. However, a separate analysis, using the set of (30) games in which both models make unique predictions for more than 90% of the possible paramenter constellations, yielded estimation results similar to those reported below.

libria; three of them have Nash Equilibria in mixed strategies, and the remaining 20 games are strict dominance solvable.

Table II contains the maximum-likelihood estimates of the two models' respective distribution parameters. The estimates were obtained by maximizing the logarithm of (2) using a standard grid-search algorithm. The maximum values of the log-likelihood functions are also given in the table, denoted by l^*. Figures 3 and 4 depict the estimated density functions of the parameter pairs $(\lambda^i; \tilde{\lambda}^j)$ and $(\epsilon^i; \tilde{\epsilon}^j)$ respectively.

Insert Table II about here.

Insert Figure 3 about here.

Insert Figure 4 about here.

As the figures illustrate, the estimations reveal considerable differences between the distributions of the subjects' actual response parameters (λ^i and ϵ^i) and the distributions of the response parameters that the subjects attribute to their opponents ($\tilde{\lambda}^j$ and $\tilde{\epsilon}^j$). In particular, it appears that subjects have a systematically distorted perception of their opponents: Both distributions of the actual response parameters λ^j and ϵ^j lie more in the "rational" area of the parameter range, as compared to the distributions of the perceived response parameters $\tilde{\lambda}^j$ and $\tilde{\epsilon}^j$. The estimated means are $E[\lambda^i] = 7.20$ as compared to $E[\tilde{\lambda}^j] = 3.92$, and $E[\epsilon^i] = 0.30$ as compared to $E[\tilde{\epsilon}^j] = 0.58$ (with estimated variances of $var(\lambda^i) = 20.83$, $var(\tilde{\lambda}^j) = 26.62$,

$\mathrm{var}(\varepsilon^i) = 0.02$, and $\mathrm{var}(\hat{\varepsilon}^j) = 0.10$).

This discrepancy between the means of the subjects' actual and perceived response parameters can be tested statistically by reestimating the models under the restriction that the expected values of the respective pairs of parameters are equal, i.e., under the null hypotheses that $E[\lambda^i] = E[\hat{\lambda}^j]$ and $E[\varepsilon^i] = E[\hat{\varepsilon}^j]$ hold respectively. Denoting the log-likelihood of the restricted model by l^r, the likelihood-ratio statistic $2(l^* - l^r)$ is (asymptotically) χ^2-distributed with one degree of freedom. The resulting critical significance levels of rejecting the null hypotheses are $p = 0.021$ for the ALE strategy model and $p = 2.749 \cdot 10^{-9}$ for the ANNE strategy model (two-tailed). Hence, the observation that subjects on average behave as if underestimating the response precision of their opponents is statistically significant in the data set, and highly so under the ANNE assumptions.[17] [18]

The related hypotheses that the actual and perceived response parameters are

[17] The observation that subjects tend to ignore their opponent's rationality has, to my knowledge, first been explicitely stated by Beard and Beil (1994) in the context of extensive-form games. See also the other experimental studies cited in Footnote 3, as well as the related papers by Goeree and Holt (2000) and Kübler and Weizsäcker (2002), all of which roughly support this hypothesis. Different structural models employing "non-equilibrium beliefs" have been estimated by Costa-Gomes and Zauner (1999) and Camerer et al. (2002).

[18] It is important to recall that the analysis in this paper only considers one-shot games. Repetitions of experimental games, or other dynamic settings, may well help the subjects learn to avoid the observed tendency to ignore the other player's rationality.

chosen from identical distributions, i.e., that both $\rho = \tilde{\rho}$ and $K = \tilde{K}$ hold under the ALE strategy model, and that both $a = \tilde{a}$ and $b = \tilde{b}$ hold under the assumptions of the ANNE strategy model, are rejected on levels of significance of $p = 5.083 \cdot 10^{-9}$ and $p = 1.227 \cdot 10^{-10}$, respectively.

Now consider the question whether the introduction of type variation, in contrast to assuming that all subject types are equal, does significantly increase the models' statistical fit in the data. More specifically, restricting both λ^i and $\tilde{\lambda}^j$ to equal a constant value, $\lambda^i = \tilde{\lambda}^j = \lambda$, amounts to assuming that all subjects play Logit Equilibrium strategies with a fixed parameter λ (see Subsection 2.1). Since this set of assumptions is also nested in the model (because the familiy of gamma distributions contains the homogeneity case, where all subjects have the same parameter), one can again perform a likelihood-ratio test, with the according test statistic being χ^2-distributed with three degrees of freedom. The maximum-likelihood estimation of λ is 3.06, and the Logit Equilibrium model is rejected on a significance level of $p = 1.083 \cdot 10^{-39}$. Analogously, the ANNE strategy model with the restriction $\varepsilon^i = \tilde{\varepsilon}^j = \varepsilon$ is rejected on a sinificance level of $p = 3.525 \cdot 10^{-63}$; with an estimated value of $\varepsilon = 0.53$.

Also, one can ask whether the hypothesis of type homogeneity can be sustained at any reasonable level of significance if one allows for a wrong perception of the opponent. That is, one may formulate the hypotheses that $\lambda^i = \bar{\lambda}$ and $\tilde{\lambda}^j = \bar{\tilde{\lambda}}$ both hold for all subjects under the ALE assumptions (for some fixed values $\bar{\lambda}$ and $\bar{\tilde{\lambda}}$), and that $\varepsilon^i = \bar{\varepsilon}$ and $\tilde{\varepsilon}^i = \bar{\tilde{\varepsilon}}$ both hold for all subjects under the ANNE assumptions (for some $\bar{\varepsilon}$ and $\bar{\tilde{\varepsilon}}$). The two according likelihood-ratio test statistics (following χ^2-distributions with two degrees of freedom) yield rejections of these sets of assumptions on the $p = 2.155 \cdot 10^{-34}$ level for the ALE strategy model, and on the $p = 2.573 \cdot 10^{-64}$ level for the ANNE strategy model. One can conclude from the tests described in this and the previous paragraph that there is substantial variation in the subject pool in terms of belief and response precision.

Nevertheless, it may be of interest to know whether the estimated perceived and actual response precisions coincide if one restricts the subject population to be homogeneous. (Under the ALE asumptions, this amounts to a test discriminating between Logit Equilibrium strategies and the more general ALE strategies.) The according ML estimates of the response parameters are $\bar{\lambda} = 4.08$ and $\bar{\tilde{\lambda}} = 2.18$ in the ALE strategy model, and $\bar{\varepsilon} = \bar{\tilde{\varepsilon}} = 0.53$ in the ANNE strategy model. Hence, the two models yield qualitatively different results under the restriction of subject homogeneity: Only the ALE estimations reveal a tendency of the subjects to underestimate the opponent's response precision (on a significance level of $p = 6.605 \cdot 10^{-8}$).

In sum, the estimations of the models using all available experimental data suggest that while any assumption of type homogeneity is rejected on very high levels of significance, there is also a significant tendency for subjects to attribute a lower degree of response precision to their opponent than they have themselves. However, one can argue that much of the evidence supporting this observation relies on the manner in which type heterogeneity was introduced. Perhaps, the assumed parameter distributions bias the estimations. To answer this question, the two models were also estimated separately for each subject, i.e., the expression (1) was maximized using each individuals's data only. The results support the observation of belief distortions made above: Under the ALE strategy model, the ML estimate of λ^i exceeds the estimate for $\tilde{\lambda}^j$ for 115 out of 160 subjects. In 50 cases was this discrepancy significant on the p = 5% level. (At p = 1%, it was significant in 14 cases, and at p = 10% in 58 cases.) The reverse relationship, $\lambda^i \leq \tilde{\lambda}^j$, was estimated to hold for 45 subjects, and for none of them does $\lambda^i < \tilde{\lambda}^j$ hold significantly at p = 5%. (For 3 subjects it holds at p = 10%.) The ANNE estimations reproduce these numbers almost exactly: $\varepsilon^i < \tilde{\varepsilon}^j$ holds at the estimated maximum of (1) in 115 cases, and significantly at p = 5% in 37 cases (at p = 1% in 18 cases, at p = 10% in 59 cases). $\varepsilon^i \geq \tilde{\varepsilon}^j$ holds in 45 cases, and again none of these show the subjects' actual error rate ε^i to be significantly larger than the perceived error rate $\tilde{\varepsilon}^j$ (not even at p = 10%). Hence, the apparent tendency to ignore the other player's response rationality is confirmed by the separate estimations for each subject.

Before moving on to results concerning subsets of data, the statistical analysis of the pooled data set is concluded by reporting goodness-of-fit tests, as well as a test of model selection between the two competing models. As a measure of the models' goodness of …t, exact chi-square tests were performed, comparing the models' predictions and the observed decisions. These tests were done for each of the 5 + 4 + 28 = 37 di¤erent player decisions in the games. (For the 14 asymmetric games by Costa-Gomes et al., 2001, row players and column players need to be considered separately.) On the level of $p = 5\%$, the predictions of the ALE strategy model (with the parameter values given in Table II) were rejected for 13 out of the 37 decision situations, and for the ANNE model in 11 out of 37 cases. Hence, while both models' predictions are accepted for the majority of the decision situations, the rejection rates are too high to lie within the limits of chance, as only 5% of the predictions should be rejected in expectation if the models were literally true. This may point at misspeci…cations of the models, or at systematic di¤erences in behavior for di¤erent subsets of games, which will be addressed below.

The fact that the ALE strategy model is rejected more often than the ANNE strategy model contrasts with the fact that the former yields a higher likelihood than the latter (see Table II). The question arises whether one model outperforms the competing model in a statistically signi…cant way. Since neither is a special case of the other, a model-selection technique for non-nested models is needed. Vuong (1989) showed that for pairs of non-nested models a surprisingly simple test can be applied, in analogy to standard likelihood-ratio tests: Let l^i_{ALE} and l^i_{ANNE} be the two models' likelihoods for the decision vector of subject i, evaluated at the ML parameter values reported in Table II. Then, clauses (i) and (iv) of Theorem 5.1 in Vuong (1989) immediately imply that the statistic

$$\frac{\sum_{i=1}^{N} \ln l^i_{ALE} - \ln l^i_{ANNE}}{\sqrt{\frac{1}{N}\sum_{i=1}^{N}(\ln l^i_{ALE} - \ln l^i_{ANNE})^2}}$$

asymptotically follows a standard normal distribution, under the null hypothesis that both models perform equally well.[19]

For the pooled data used here, the test shows that the ALE strategy model outperforms the ANNE strategy model significantly, on a level of $p = 4.553\%$ (two-tailed). The next subsection will demonstrate, however, that this result does not hold for all subsets of the data.

3.3 Robustness: Estimation results using subsets of data

To give an impression of the robustness of the results, the data are separated following two criteria. First, the models are estimated for each of the three experiments, in order to control for possible treatment effects. Second, the data are divided according to systematic differences between the games played.

[19] For this result to apply it is essential (as it is for the other likelihood-ratio tests used in this paper) that the subjects' types are drawn independently. The numerator of the test statistic is simply the likelihood ratio of the two competing models, and the denominator is a consistent estimator of this ratio's variance, under the unknown true distribution of the data generating process.

Tables III and IV, in Appendix C, contain the estimation results of the two models if the data from the three experiments are used separately. Also, the tables contain: The significance levels of rejecting four hypotheses analogous to those that were tested in the previous subsection; the number of different decision situations for which the models' predictions are rejected by exact chi-square tests on the level of $p = 5\%$; and – in the lower panels of the tables – the number of subjects for which the individual estimations yield each of the two possible parameter asymmetries (with cases significant at $p = 5\%$ in parentheses). As the tables show, the results are qualitatively very similar to the results of the pooled data reported above, at least for two of the three data sets. In both the Stahl and Wilson (1994) data and the data by Costa-Gomes et al. (2001), the estimated actual response functions of both models are on average closer to the best response function than the estimated perceived response funtions are. I.e., the subjects tend to ignore their opponents' rationality, according to these estimation results (see the estimates of $E[\lambda^i]$, $E[\tilde{\lambda}^j]$, $E[\varepsilon^i]$, and $E[\tilde{\varepsilon}^j]$, reported in the tables). These differences are, as indicated in the first rows of the tables' second sections, highly significant in both of these data sets and in both models, all corresponding levels of significance being below $p = 10^{-6}$. Furthermore, as in the analysis of the pooled data, all restrictions of type heterogeneity among the subjects are rejected on even higher levels of significance.

For the remaining data set by Stahl and Wilson (1995), the results of the previous subsection are only partly reproduced. While any assumption of type homogeneity is again strongly rejected, the result of underestimated precision levels shows up only in the ANNE strategy model, with a signi…cance of p = 0:026. In the ALE strategy model, the subjects' perceived response precision is, on average, estimated to be higher than the actual response precision, on a level of signi…cance of p = 0:062. Apparently, the assumptions concerning the error structure make a signi…cant di¤erence here. However, a possible reason for this deviation from the general result pattern is that the ALE strategy model is poorly identi…ed in the Stahl and Wilson (1995) data set (which is the smallest data set used). This is indicated by the model selection test by Vuong (1989): In the Stahl and Wilson (1995) data, ANNE strongly outperforms ALE according to this test, on a signi…cance level of p = 0:001. In the Stahl and Wilson (1994) data, ANNE insigni…cantly outperforms ALE, at p = 0:219. In the data by Costa-Gomes et al. (2001), ALE outperforms ANNE at p = 0:009.[20]

Another question is whether distinctive properties of some of the games lead subjects to exhibit a choice behavior of the speci…c kind described above. In particular, it is possible that subjects simply fail to identify dominance relations among the opponent's strategies, but never play a dominated strategy themselves.[21] Since most of

[20] A comparison of the two sets of experimental instructions by Stahl and Wilson (1994, 1995) does not reveal strong di¤erences which could account for large variation in behavior.

[21] This pattern, which can be seen as re‡ecting the ignorance concerning the opponent's rationality in cases of games with dominated strategies, is suggested almost immediately by an inspection of

the games used contain dominated strategies, such a failure may drive the results of the analysis using the pooled data. To investigate this possibility, the data analysis is repeated for two partitions of the set of games: For the first partition, define the set DT as the set containing all games in which at least one player has a dominant strategy. (All of these 10 games are from Costa-Gomes et al., 2001.) This set is compared with NDT, the set of games without dominant strategies. Subsequently, a similar comparison is made between the set DD, containing all games with dominated strategies (for either player), and NDD, the set of games without dominated strategies.[22]

Tables V and VI show the results of the estimations for the DT and the NDT data, organized as in the preceeding tables. In both data sets, the estimated average

---

the aggregate choice data in Appendix B.

[22] More precisely, the set NDD contains all games in which no strategy is dominated by a pure strategy, and DD contains the remaining games. Unfortunately, the set of games in which no strategy is dominated by any (pure or mixed) strategy contains only the three games with mixed strategy Nash Equilibria, so I decided to include the four games with strategies that are solely dominated by mixed strategies into NDD, for two reasons: First, it increases the amount of data in the NDD analysis. Second, it mitigates the problem that the results are subject to any systematic differences in behavior appearing only in games without pure Nash Equilibria that are not captured by the models. While the estimations from the NDD data will certainly have to be interpreted with caution, it also turned out in a separate analysis, using the set containing only the three games with mixed Nash Equilibria, that the results are qualitatively similar (see below, Footnote 24).

response precisions are, once again, di¤erent between the actual and the perceived response functions of both models, with a tendency to underestimate the opponents' response precision (though on lower levels of signi…cance than in the pooled data). Also, both models' special cases with the restriction to type homogeneity are strongly rejected, in both data sets. Hence, it appears that the distinction between games with and without dominant strategies leaves the main results of the previous analysis essentially untouched.

For the sets of games with and without dominated strategies, DD and NDD, the results are summarized in Tables VII and VIII. Here, the two data sets yield somewhat di¤erent results concerning the absolute and relative estimated levels of the average response parameters: While the games with dominated strategies induce the familiar pattern of subjects underestimating their opponents' response precision, the tables show that for the NDD data set this behavior appears only under the assumptions of the ANNE strategy model, whereas the ALE estimations result in insigni…cant di¤erences between the mean levels of λ^i and $\tilde{\lambda}^j$. Therefore, the evidence for the behavioral pattern described above is weaker for these games.

However, the test by Vuong (1989) indicates that ANNE outperforms ALE in the NDD data, at a level of $p = 6.914 \cdot 10^{-17}$. In the other data subsets, the model selection test does not yield strong rejections of either model. In the DT and DD data, ALE outperforms ANNE at signi…cance levels of $p = 0.102$ and $p = 0.126$, respectively. In the NDT data, ANNE outperforms ALE at $p = 0.109$. Also, note that under both models, the estimated actual response precisions are much lower in the NDD data than in the DD data (so fewer decisions are estimated to be made randomly in the DD data), pointing at a behavior more consistent with the model predictions in the "simpler" games contained in DD.[23] [24]

This discrepancy is confirmed by the results of the goodness-of-fit tests, reported in the third section of each of the tables, as the model rejection rate is smaller for the DD data than it is for the NDD data. An analogous difference appears between the DT data and the NDT data. The goodness-of-fit tests also show that the rejection rates are only slightly reduced when subsets of data are considered separately, as compared to the pooled-data analysis of the previous subsection.

In sum, the separation of the data into subsets mostly confirms the observed tendency to ignore or underestimate the opponent's response rationality. The corresponding effects are statistically significant in all but two cases of (model, data set)

[23] The lower response precision in NDD may partly explain the apparent identification problem of the ALE strategy model in these data: With a high error rate (low λ^i) little can be inferred about the subjects' beliefs. It is not clear, however, why such problems should matter less in estimations of the ANNE strategy model.

[24] The estimated mean values of the models' parameters using only the set of games with mixed equilibria (which contains three of the seven games in NDD, cf. Footnote 22) are $E[\lambda^i] = 5.39$; $E[\tilde{\lambda}^j] = 7.23$, $E[\varepsilon^i] = 0.62$; and $E[\tilde{\varepsilon}^j] = 0.90$. The significance levels of rejecting the hypotheses $E[\lambda^i] = E[\tilde{\lambda}^j]$ and $E[\varepsilon^i] = E[\tilde{\varepsilon}^j]$ are $p = 0.661$ and $p = 0.035$, respectively.

combinations. In these two cases, the model used is strongly rejected in favor of the competing model, which reproces the above observation in both cases, using the same data.

4 Conclusions

In this paper, two behavioral model of beliefs and responses are presented and estimated from experimental data. The estimation results show that while there is a large type variation among the subjects in each of the experiments, the subjects on average are prone to make a systematic prediction error: Under the assumptions of the models, subjects act as if underestimating the response precision of their opponents.

A rather conservative interpretation of this result is to view it as a critique of an assumption which is typically made in game-theoretic models of quantal response: In these models, one usually supposes that experimental subjects are aware of the level of randomness in their opponents' motivations (although the experimenter herself is not). This assumption is relaxed in the above estimations, and the corresponding statistical tests indicate that it is consistently rejected in the data.

Taking a broader view, one may read the estimation results as an indication of a more general "anomaly": Subjects tend to ignore their opponents' incentives. One can then ask whether this evidence is an artifact of the experimental environments in the laboratories, and in particular of the fact that only normal-form games with abstract matrix presentations were used in the analysis. Perhaps, adding a context to the experiments would help the subjects to see their opponent's decision problems more vividly and clearer.

References

Beard, T.R., and Beil, R.O. (1994). "Do people rely on the self-interested maximization of others? An experimental test," Management Science 40, 252-262.

Camerer, C., Ho, T.H., and Chong, J.-K. (2002). "Sophisticated Experience-Weighted Attraction learning and teaching in repeated games," Journal of Economic Theory 104, 137-188.

Costa-Gomes, M., and Zauner, K.(1999). "Learning, non-equilibrium beliefs, and non-pecuniary payo¤ uncertainty in an experimental game," Economic Theory, forthcoming.

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

Goeree, J.K., and Holt, C.A. (2000). "A model of noisy introspection," mimeo, University of Virginia.

Haruvy, E., and Stahl, D.O. (1998). "An empirical model of equilibrium selection in symmetric normal-form games," Working Paper 9803, Department of Eonomics, University of Texas at Austin.

Ho, T.H., Camerer, C., and Weigelt, K. (1998). "Iterated dominance and iterated best response in experimental 'p-beauty contests'," American Economic Review 88, 947-969.

Huck, S., and Weizsäcker, G. (2002). "Do players correctly estimate what others do? Evidence of conservatism in beliefs," Journal of Economic Behavior and Organization 47, 71-85.

Kübler, D., and Weizsäcker, G. (2002). "Limited depth of reasoning and failure of cascade formation in the laboratory," mimeo, Humboldt University Berlin.

McKelvey, R.D., and Palfrey, T.R. (1995). "Quantal response equilibria for normal form games," Games and Economic Behavior 10, 6-38.

McKelvey, R.D., and Palfrey, T.R. (1998). "Quantal response equilibria for extensive form games," Experimental Economics 1, 9-41.

McKelvey, R.D., Turocy, T., and Grayver, E. (1996). "Gambit – Graphics user interface," http://www.hss.caltech.edu/~gambit/Gambit.html.

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

Schotter, A., Weigelt, K., and Wilson, C. (1994). "A laboratory investigation of mutiperson rationality and presentation e¤ects," Games and Economic Behavior 6, 445-468.

Stahl, D.O. (1999). "Evidence based rules and learning in symmetric normal-form games," International Journal of Game Theory 28, 111-130.

Stahl, D.O., and Wilson, P.W. (1994). "Experimental evidence on players' models of other players," Journal of Economic Behavior and Organization 25, 309-327.

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

Van Huyck, J.B., Battalio, R.C., and Beil, R.O. (1990). "Tacit Coordination games, strategic uncertainty, and coordination failure," American Economic Review 80, 234-248.

Van Huyck, J.B., Battalio, R.C., and Beil, R.O. (1991). "Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games," Quarterly Journal of Economics 106, 885-910.

Vuong, Q.H. (1989). "Likelihood ratio tests for model selection and non-nested hypotheses," Econometrica 57, 307-333.

Zauner, K. (1999). "A payo¤ uncertainty explanation of results in experimental centipede games," Games and Economic Behavior 26, 157-185.

Appendix

A A random-utility interpretation of ANNE strategies

In the following, for any given normal-form game a corresponding game with incomplete information is constructed, which has a Bayesian Nash Equilibrium strategy that is observationally equivalent to an ANNE strategy of the original game. Thereby, existence of ANNE strategies is established.

Summarized in words, the original game is modified in order to make uniform "trembles" an optimal strategy in the modified game. From the view of an outside observer, player a will put probability weight ε^a on the uniform distribution, and will put the remaining mass on her best response to player b's strategy. Also, the information structure is chosen such that, from player a's perspective, player b adopts an analogous behavior (with weight $\hat{\varepsilon}^b$ on the uniform distribution) regardless of player a's private information. The realization of nature's initial draw in the modified game corresponds to the random choices of the players.

For a given game Γ and any pair $(\varepsilon^a; \hat{\varepsilon}^b) \in [0;1] \times [0;1]$, let the information structure of the corresponding game of incomplete information be given by

$$I_{\Gamma, \varepsilon^a, \hat{\varepsilon}^b} = \{\Omega; P^a; P^b; \mu\}, \tag{3}$$

where

$$\Omega = \{A; B_1; ...; B_{m^b}; C_1; ...; C_{m^a}; D_{1;1}; D_{1;2}; ...; D_{m^a;m^b}\} \tag{4}$$

is the set of possible states of nature of the game, and

$$
\begin{aligned}
P^a &= \{\{A; B_1; \ldots; B_{m^b}\}; \{C_1; D_{1;1}; \ldots; D_{1;m^b}\}; \ldots; \{C_{m^a}; D_{m^a;1}; \ldots; D_{m^a;m^b}\}\}; \\
P^b &= \{\{A; C_1; \ldots; C_{m^a}\}; \{B_1; D_{1;1}; \ldots; D_{m^a;1}\}; \ldots; \{B_{m^b}; D_{1;m^b}; \ldots; D_{m^a;m^b}\}\}
\end{aligned}
\tag{5}
$$

are the partitions of Ω of the two players. For example, if the true state is $\omega = B_1$, a knows that $\omega \in \{A; B_1; \ldots; B_{m^b}\}$, and b knows that $\omega \in \{B_1; D_{1;1}; \ldots; D_{m^a;1}\}$. Using the payoff manipulations introduced below, this corresponds to the case that player a views the game as the original game (has the same payoff as in state A), but b's payoffs are modified such that he optimally chooses his first strategy. He is, however, unsure whether a also has modified payoffs $(D_{1;1}; \ldots; D_{m^a;1})$ or not. Denote by $P^i(\omega)$ player i's information set if ω is drawn. Furthermore, let

$$
\begin{aligned}
\mu = {} & ((1-\varepsilon^a)(1-\tilde{\varepsilon}^b); \frac{\tilde{\varepsilon}^b(1-\varepsilon^a)}{m^b}; \ldots; \frac{\tilde{\varepsilon}^b(1-\varepsilon^a)}{m^b}; \frac{\varepsilon^a(1-\tilde{\varepsilon}^b)}{m^a}; \ldots; \frac{\varepsilon^a(1-\tilde{\varepsilon}^b)}{m^a}; \\
& \frac{\varepsilon^a\tilde{\varepsilon}^b}{m^a m^b}; \ldots; \frac{\varepsilon^a\tilde{\varepsilon}^b}{m^a m^b})
\end{aligned}
\tag{6}
$$

be the common prior over Ω, where all elements $B_l;\ l = 1; \ldots; m^b$, occur with probability $\frac{\tilde{\varepsilon}^b(1-\varepsilon^a)}{m^b}$, all $C_k; k = 1; \ldots; m^a$, with probability $\frac{\varepsilon^a(1-\tilde{\varepsilon}^b)}{m^a}$, and all $D_{k;l}; k = 1; \ldots; m^a; l = 1; \ldots; m^b$, with probability $\frac{\varepsilon^a\tilde{\varepsilon}^b}{m^a m^b}$.

Given this information structure, a strategy for player i is a mapping $f^i : \Omega \to Q^i$ where $f^i(\cdot)$ satisfies the condition that $P^i(\omega') = P^i(\omega)$ implies $f^i(\omega') = f^i(\omega)$. Let S^i be the set of all f^i satisfying this condition, i.e., S^i is player i's strategy set.

(Importantly, the sets of both players' mixed action profiles, Q^a and Q^b, remain unchanged between the original and the modified game.)

Now consider the following manipulations of the players' utility functions: The state-dependent utility function of player a, $u^a_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; \omega)$, remains unchanged as compared to a's utility in the original game if ω lies in $\{A; B_1; \ldots; B_{m^b}\}$, i.e., $\forall l = 1; \ldots; m^b : u^a_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; A) = u^a_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; B_l) = u^a(a^a; a^b)$. In the case $\omega \in \{C_k; D_{k;1}; \ldots; D_{k;m^b}\}$, a's utility is modified such that a chooses optimally the kth action from her action set: $\forall k = 1; \ldots; m^a; l = 1; \ldots; m^b$;

$$u^a_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; C_k) = u^a_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; D_{k;l}) = \begin{Bmatrix} u^a(a^a; a^b) & \text{if } a^a = a^a_k \\ -F & \text{otherwise} \end{Bmatrix}; \tag{7}$$

where F satisfies $-F < \min_{i;a^a;a^b} u^i(a^a; a^b)$. Analogously for player b, $\forall k = 1; \ldots; m^a$; $l = 1; \ldots; m^b : u^b_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; A) = u^b_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; C_k) = u^b(a^a; a^b)$, and

$$u^b_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; B_l) = u^b_{\varepsilon^a,\tilde{\varepsilon}^b}(a^a; a^b; D_{k;l}) = \begin{Bmatrix} u^b(a^a; a^b) & \text{if } a^b = a^b_l \\ -F & \text{otherwise} \end{Bmatrix}: \tag{8}$$

Definition 4 $\Gamma_{\varepsilon^a,\tilde{\varepsilon}^b}$ is the $(\varepsilon^a; \tilde{\varepsilon}^b)$-perturbed version of Γ if $\Gamma_{\varepsilon^a,\tilde{\varepsilon}^b} = \{a; b; I_{\Gamma_{\varepsilon^a,\tilde{\varepsilon}^b}}; A^a; A^b; u^a_{\varepsilon^a,\tilde{\varepsilon}^b}; u^b_{\varepsilon^a,\tilde{\varepsilon}^b}\}$; where A^a and A^b are defined as in Γ and $I_{\Gamma_{\varepsilon^a,\tilde{\varepsilon}^b}}$; $u^a_{\varepsilon^a,\tilde{\varepsilon}^b}$; and $u^b_{\varepsilon^a,\tilde{\varepsilon}^b}$ are given by expressions (3) to (8).

Proposition 5 The following are equivalent: (i) The player-a strategy π^a is an ANNE strategy with parameters $(\varepsilon^a; \tilde{\varepsilon}^b)$ of game Γ. (ii) π^a is player a's probability distribution vector resulting from a Bayesian Nash Equilibrium strategy of game $\Gamma_{\varepsilon^a,\tilde{\varepsilon}^b}$.

Proof. In any Bayesian Nash Equilibrium of $\Gamma_{\varepsilon^a,\varepsilon^b}$, the optimality condition for player a implies that if it holds for some k ($k = 1, \ldots, m^a$) that $\omega \in \{C_k, D_{k,1}, \ldots, D_{k,m^b}\}$ – which occurs with probability $\frac{\varepsilon^a}{m^a}$ –, then player a chooses action a^a_k. Analogously for player b, if $\omega \in \{B_l, D_{1,l}, \ldots, D_{m^a,l}\}$, then a^b_l is chosen, $\forall l = 1, \ldots, m^b$. Adding up the corresponding probability weights and using the uniform functions P^i_0 defined in Section 2, one can hence write any Bayesian Nash Equilibrium distribution (π^a, π^b) as $\pi^a = (1 - \varepsilon^a)\hat{\pi}^a + \varepsilon^a P^a_0$ and $\pi^b = (1 - \varepsilon^b)\hat{\pi}^b + \varepsilon^b P^b_0$, for some strategy pair $(\hat{\pi}^a, \hat{\pi}^b)$. More specifically, $\hat{\pi}^a$ is a's optimal probability distribution over A^a in the case $\omega \in \{A, B_1, \ldots, B_{m^b}\}$, and $\hat{\pi}^b$ is optimal for b in the case $\omega \in \{A, C_1, \ldots, C_{m^a}\}$. Notice that for any information set of a it holds that a's updated probability of the event $\omega \in \{B_l, D_{1,l}, \ldots, D_{m^a,l}\}$ is equal to $\frac{\varepsilon^b}{m^b}$, $\forall l = 1, \ldots, m^b$, and that a's updated probability of the event $\omega \in \{A, C_1, \ldots, C_{m^a}\}$ is $(1 - \varepsilon^b)$, so in equilibrium player a does not update her expectation about b's information set (equivalently, about b's behavior) after receiving her private information. In particular, if $\omega \in \{A, B_1, \ldots, B_{m^b}\}$, a expects b to play according to the distribution $\pi^b = (1 - \varepsilon^b)\hat{\pi}^b + \varepsilon^b P^b_0$ – as stated above – and hence chooses $\hat{\pi}^a$ so as to solve

$$\max_{\pi \in A^a} (1 - \varepsilon^b) u^a(\pi, \hat{\pi}^b) + \varepsilon^b u^a(\pi, P^b_0),$$

which is exactly a's maximization problem in calculating an ANNE strategy for parameters $(\varepsilon^a, \varepsilon^b)$ in game Γ. This last statement is true because if $\omega \in \{A, B_1, \ldots, B_{m^b}\}$ then a's utility is the same as in Γ. An analogous observation applies to player b, so the two …xed-point problems are equivalent. ■

Using this reinterpretation of an ANNE strategy as a Bayesian Nash Equilibrium strategy, existence of ANNE strategies for any $({}^{2}{}^{a}; ê^{b})$ 2 $[0;1]$ £ $[0;1]$ follows from Nash's Theorem. Also, it holds by the de…nition of trembling-hand perfection that any Nash Equilibrium of a game which is the limit of a sequence of ANNE strategies with parameters $({}^{2}{}^{i}; ê^{j})$ approaching $(0;0)$ is necessarily trembling-hand perfect.

B Games

Figures 5, 6, and 7 depict the 23 games used in the analysis of Section 3, taken from Stahl and Wilson (1994, 1995) and Costa-Gomes et al. (2001). The …gures show the games using the point numbers that were presented to the subjects. Stahl and Wilson (1994) determined a subject's earnings for a given game by calculating the subject's expected outcome from playing against the distribution generated by pooling the choices of all her possible opponents. The resulting number of points was then used as the subject's number of winning chips (out of 100) in a binary-lottery for winning a …xed prize of $2.50. In Stahl and Wilson (1995), an analogous procedure was applied using a prize of $2.00 for each game. Costa-Gomes et al. (2001) paid o¤ one of the 18 games, paying 40 cents for each point as given in the tables. 36 of the 72 subjects were assigned the role of the column player. In the experiment, the games were presented to them as if they were row players in the "transposed" games, to guarantee equal conditions between the two groups.

Insert Figure 5 about here.

Insert Figure 6 about here.

Insert Figure 7 about here.

C Tables

Insert Table III about here.

Insert Table IV about here.

Insert Table V about here.

Insert Table VI about here.

Insert Table VII about here.

Insert Table VIII about here.

Table I: Game $¡_1$

	L	R
U	75; 51	42; 27
D	48; 80	89; 68

Table II: ML estimates of distribution parameters, using pooled experimental data.

ALE parameters		ANNE parameters	
ρ :	0.35	a :	2.28
K :	2.49	b :	5.28
$\tilde{\rho}$:	0.15	$\tilde{a}$:	0.77
$\tilde{K}$:	0.58	$\tilde{b}$:	0.55
l^*	-907.154	l^*	-930.039

Table III: Separate estimations of the ALE strategy model for the data from Stahl and Wilson (1994, 1995) and Costa-Gomes et al. (2001).

Data set:	SW 1994	SW 1995	CGCB 2001
ρ	6.55	7.21	0.23
κ	171.24	53.06	1.79
$\tilde{\rho}$	0.43	0.07	2.08
$\tilde{\kappa}$	0.75	0.82	4.60
l^*	-140.468	-151.128	-571.853
$E[\lambda^i]$	26.13	7.36	7.74
$E[\tilde{\lambda}^j]$	1.75	11.33	2.21
sig($E[\lambda^i] = E[\tilde{\lambda}^j]$)	3.322×10^{-10}	0.062	2.078×10^{-7}
sig($\rho = \tilde{\rho}; \kappa = \tilde{\kappa}$)	8.623×10^{-13}	0.019	8.620×10^{-8}
sig($(\bar{\lambda}^i, \bar{\tilde{\lambda}}^j)$ fix over i)	4.706×10^{-11}	3.443×10^{-6}	2.258×10^{-25}
sig($\bar{\lambda}^i = \bar{\tilde{\lambda}}^j = \lambda$, fix over i)	4.720×10^{-19}	1.159×10^{-5}	7.509×10^{-31}
# of model rejections at p = 5%	0 (out of 5)	1 (out of 4)	8 (out of 28)
# of subjects with $\lambda^i > \tilde{\lambda}^j$	29 (16 sig.)	25 (9 sig.)	61 (26 sig.)
# of subjects with $\lambda^i \leq \tilde{\lambda}^j$	11 (0 sig.)	23 (0 sig.)	11 (0 sig.)

Note: Significance level for the significant cases in the last two rows is 5%.

Table IV: Separate estimations of the ANNE strategy model for the data from Stahl and Wilson (1994, 1995) and Costa-Gomes et al. (2001).

Data set:	SW 1994	SW 1995	CGCB 2001
a	10.53	11.20	2.43
b	160.51	36.19	4.38
$\hat{a}$	0.91	0.21	3.99
$\hat{b}$	0.57	0.36	2.40
l^*	-134.444	-142.432	-622.046
$E[\varepsilon^i]$	0.06	0.24	0.36
$E[\hat{\varepsilon}^j]$	0.61	0.37	0.62
$sig(E[\varepsilon^i] = E[\hat{\varepsilon}^j])$	2.936×10^{-8}	0.026	2.473×10^{-7}
$sig(a = \hat{a}; b = \hat{b})$	9.590×10^{-8}	0.006	1.658×10^{-6}
$sig((\bar{\varepsilon}^i; \bar{\hat{\varepsilon}}^j)$ fix over $i)$	2.356×10^{-23}	3.379×10^{-11}	5.409×10^{-34}
$sig(\bar{\varepsilon}^i = \bar{\hat{\varepsilon}}^j = \varepsilon;$ fix over $i)$	2.532×10^{-24}	1.910×10^{-10}	5.376×10^{-33}
# of model rejections at $p = 5\%$	0 (out of 5)	0 (out of 4)	10 (out of 28)
# of subjects with $\varepsilon^i < \hat{\varepsilon}^j$	28 (14 sig.)	25 (9 sig.)	61 (11 sig.)
# of subjects with $\varepsilon^i \geq \hat{\varepsilon}^j$	12 (0 sig.)	23 (0 sig.)	11 (0 sig.)

Note: See Table III.

Table V: Separate estimations of the ALE strategy model for the data from games with and without dominant strategies, DT and NDT, respectively.

Data set:	DT	NDT
ρ	0.09	0.52
κ	1.17	2.93
$\tilde{\rho}$	0.53	0.17
$\tilde{\kappa}$	1.69	0.59
l^*	-338.142	-562.328
$E[\lambda^i]$	12.73	5.63
$E[\tilde{\lambda}^j]$	3.19	3.42
$sig(E[\lambda^i] = E[\tilde{\lambda}^j])$	0.013	0.072
$sig(\rho = \tilde{\rho}; \kappa = \tilde{\kappa})$	0.001	$1.003 \cdot 10^{-4}$
$sig((\bar{\lambda}^i; \bar{\tilde{\lambda}}^j)$ fix over i)	$1.852 \cdot 10^{-28}$	$2.277 \cdot 10^{-7}$
$sig(\bar{\lambda}^i = \bar{\tilde{\lambda}}^j = \lambda$; fix over i)	$8.037 \cdot 10^{-33}$	$1.742 \cdot 10^{-8}$
# of model rejections at p = 5%	5 (out of 20)	10 (out of 17)
# of cases with $\lambda^i > \tilde{\lambda}^j$	62 (15 sig.)	96 (39 sig.)
# of cases with $\lambda^i \leq \tilde{\lambda}^j$	10 (0 sig.)	64 (0 sig.)

Note: See Table III.

Table VI: Separate estimations of the ANNE strategy model for the data from games with and without dominant strategies, DT and NDT, respectively.

Data set:	DT	NDT
a	0.86	6.63
b	2.56	12.22
$\tilde{a}$	0.37	0.93
$\tilde{b}$	0.45	0.58
l^*	-347.113	-555.608
$E[\epsilon^i]$	0.25	0.35
$E[\tilde{\epsilon}^j]$	0.45	0.62
$sig(E[\epsilon^i] = E[\tilde{\epsilon}^j])$	0.023	$9.917 \cdot 10^{-4}$
$sig(a = \tilde{a}; b = \tilde{b})$	0.008	0.004
$sig((\epsilon^i; \bar{\tilde{\epsilon}}^j)$ fix over $i)$	$4.203 \cdot 10^{-34}$	$1.423 \cdot 10^{-22}$
$sig(\epsilon^i = \bar{\tilde{\epsilon}}^j = \epsilon;$ fix over $i)$	$4.185 \cdot 10^{-33}$	$5.237 \cdot 10^{-27}$
# of model rejections at $p = 5\%$	3 (out of 20)	6 (out of 17)
# of cases with $\epsilon^i < \tilde{\epsilon}^j$	30 (15 sig.)	96 (39 sig.)
# of cases with $\epsilon^i \geq \tilde{\epsilon}^j$	42 (0 sig.)	64 (0 sig.)

Note: See Table III.

Table VII: Separate estimations of the ALE strategy model for the data from games with and without dominated strategies, DD and NDD, respectively.

Data set:	DD	NDD
ρ	0.09	38.70
κ	1.27	171.03
$\tilde{\rho}$	1.43	8.04
$\tilde{\kappa}$	3.03	35.94
l^*	-583.582	-326.213
$E[\lambda^i]$	13.14	4.42
$E[\tilde{\lambda}^j]$	2.12	4.47
$sig(E[\lambda^i] = E[\tilde{\lambda}^j])$	$5.045 \cdot 10^{-9}$	0.942
$sig(\rho = \tilde{\rho}; \kappa = \tilde{\kappa})$	$4.647 \cdot 10^{-12}$	0.976
$sig((\bar{\lambda}^i, \bar{\tilde{\lambda}}^j)\ \forall \text{ over } i)$	$2.924 \cdot 10^{-29}$	0.976
$sig(\bar{\lambda}^i = \bar{\tilde{\lambda}}^j = \lambda, \forall \text{ over } i)$	$2.073 \cdot 10^{-36}$	0.997
# of model rejections at $p = 5\%$	7 (out of 28)	4 (out of 9)
# of cases with $\lambda^i > \tilde{\lambda}^j$	106 (29 sig.)	79 (9 sig.)
# of cases with $\lambda^i \le \tilde{\lambda}^j$	54 (0 sig.)	81 (0 sig.)

Note: See Table III.

Table VIII: Separate estimations of the ANNE strategy model for the data from games with and without dominated strategies, DD and NDD, respectively.

Data set:	DD	NDD
a	0.86	9.71
b	2.64	14.71
$\tilde{a}$	1.35	57.62
$\tilde{b}$	1.18	20.72
l^*	-596.819	-316.254
$E[\epsilon^i]$	0.25	0.40
$E[\tilde{\epsilon}^j]$	0.53	0.73
$sig(E[\epsilon^i] = E[\tilde{\epsilon}^j])$	$5.997 \cdot 10^{-9}$	$3.778 \cdot 10^{-7}$
$sig(a = \tilde{a}; b = \tilde{b})$	$4.495 \cdot 10^{-8}$	$2.578 \cdot 10^{-6}$
$sig((\epsilon^i; \bar{\tilde{\epsilon}}^j)$ fix over i)	$1.712 \cdot 10^{-40}$	0.001
$sig(\epsilon^i = \bar{\tilde{\epsilon}}^j = \epsilon$; fix over i)	$1.312 \cdot 10^{-39}$	$1.592 \cdot 10^{-7}$
# of model rejections at $p = 5\%$	6 (out of 28)	3 (out of 9)
# of cases with $\epsilon^i < \tilde{\epsilon}^j$	104 (32 sig.)	108 (9 sig.)
# of cases with $\epsilon^i \geq \tilde{\epsilon}^j$	56 (0 sig.)	52 (0 sig.)

Note: See Table III.

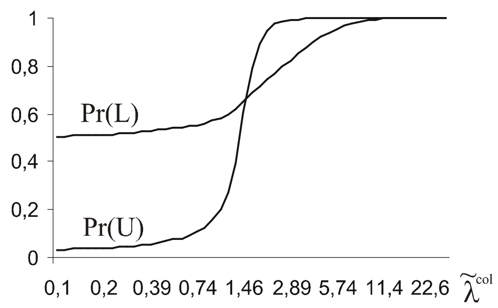

Figure 1a.

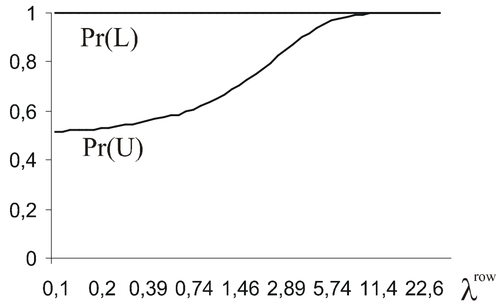

Figure 1b.

Figure 1: ALE strategies for game Γ_1. Figure 1a depicts the values of Pr(U) and Pr(L) for $\lambda^{row} = 16$ (or $\tilde{\lambda}^{row} = 16$, from the perspective of the column player), Figure 1b depicts Pr(U) and Pr(L) for $\tilde{\lambda}^{col} = 16$ (or $\lambda^{col} = 16$).

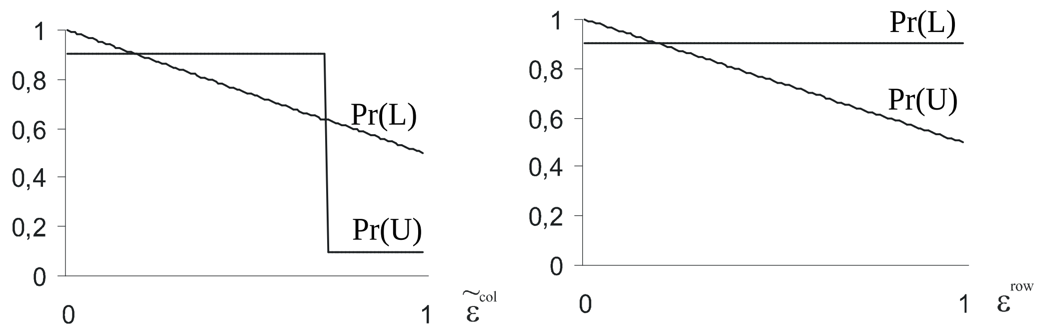

Figure 2a. Figure 2b.

Figure 2: ANNE strategies of game Γ_1. Figure 2a depicts the values of Pr(U) and Pr(L) for $\varepsilon^{row} = 0.2$ (or $\tilde{\varepsilon}^{row} = 0.2$, from the perspective of the column player), Figure 2b depicts Pr(U) and Pr(L) for $\tilde{\varepsilon}^{col} = 0.2$ (or $\varepsilon^{col} = 0.2$).

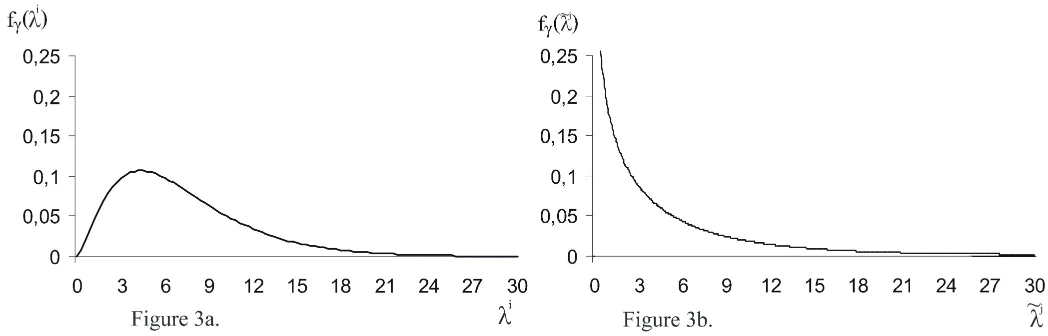

Figure 3: Estimated densities of ALE parameters using pooloed data. Figure 3a depicts the estimated density $f_\gamma(\lambda^{i}|\rho; K)$, Figure 3b the estimated density $f_\gamma(\tilde{\lambda}^{j}|\tilde{\rho}; \tilde{K})$.

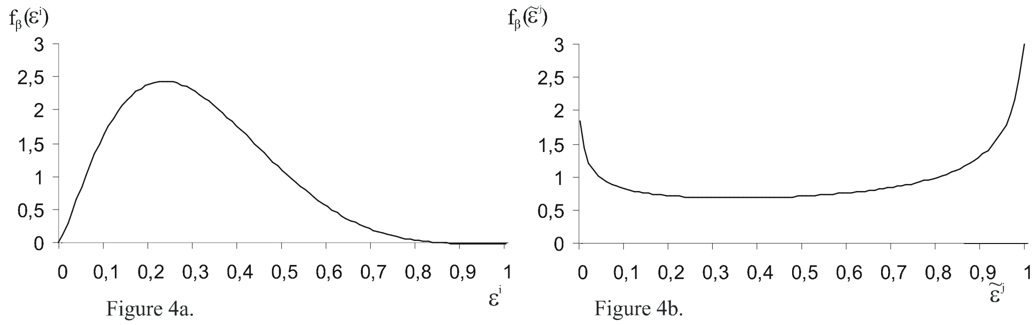

Figure 4: Estimated densities of ANNE parameters using pooloed data. Figure 4a depicts the estimated density $f_\beta(\varepsilon^i|\alpha;\beta)$, Figure 4b the estimated density $f_\beta(\tilde{\varepsilon}^j|\tilde{\alpha};\tilde{\beta})$.

Game 1:

		'T'	'M'	'B'
11	T	40;40	10;20	70;30
0	M	20;10	80;80	0;100
29	B	30;70	100;0	60;60

Game 2:

		'T'	'M'	'B'
26	T	20;20	0;60	100;0
7	M	60;0	20;20	0;60
7	B	0;100	60;0	40;40

Game 3:

		'T'	'M'	'B'
14	T	20;20	30;40	100;30
26	M	40;30	40;40	60;0
0	B	30;100	0;60	40;40

Game 5:

		'T'	'M'	'B'
18	T	10;10	100;0	20;20
0	M	0;100	70;70	30;50
22	B	20;20	50;30	40;40

Game 8:

		'T'	'M'	'B'
8	T	10;10	100;0	0;78
13	M	0;100	60;60	70;40
19	B	78;0	40;70	20;20

Figure 5: Games from Stahl and Wilson (1994).

Game 1:

		'T'	'M'	'B'
7	T	25;25	30;40	100;31
40	M	40;30	45;45	65;0
1	B	31;100	0;65	40;40

Game 4:

		'T'	'M'	'B'
26	T	30;30	50;40	100;35
15	M	40;50	45;45	10;60
7	B	35;100	60;10	0;0

Game 5:

		'T'	'M'	'B'
14	T	10;10	100;0	40;20
3	M	0;100	70;70	50;50
31	B	20;40	50;50	60;60

Game 12:

		'T'	'M'	'B'
26	T	40;40	15;22	70;30
3	M	22;15	80;80	0;100
19	B	30;70	100;0	55;55

Figure 6: Games from Stahl and Wilson (1995).

Game 2A:

		21 L	15 R
27	U	·55;79	84;52
9	D	31;46	72;93

Game 2B:

		26 L	10 R
33	U	·38;57	94;23
3	D	14:18	45;89

Game 3A:

		33 L	3 R
25	U	·75;51	42;27
11	D	48;80	89;68

Game 3B:

		34 L	2 R
26	U	55;36	16;12
10	D	21;92	87;43

Game 4A:

		31 L	0 M	5 R
25	T	70;52	38;29	37;23
11	B	46;83	59;58	85;61

Game 4B:

		32 L	4 R
14	T	68;46	31;32
2	M	47;61	72;43
19	B	43;84	91;65

Game 4C:

		20 L	16 R
33	T	51;69	82;45
0	M	28;37	57;58
3	B	22;36	60;84

Game 4D:

		28 L	1 M	7 R
32	T	42;64	57;43	80;39
4	B	28,27	39;68	61;87

Game 5B:

		25 L	11 R
4	T	74;62	43;40
5	M	25;12	76;93
27	B	59;37	94;16

Game 6B:

		6 L	4 M	26 R
22	T	64;76	14;27	39;61
14	B	42;45	95;78	18;96

Game 7B:

		7 L	2 M	27 R
20	T	56;78	23;53	89;49
16	B	31;35	95;64	67;91

Game 8B:

		17 L	19 R
8	T	71;49	28;24
3	M	46;16	57;88
25	B	42;82	84;60

Game 9A:

		23 L	13 R
33	T	45;66	82;31
0	TM	22;14	57;55
0	BM	30;42	28;37
3	B	15:60	61;88

Game 9B:

		31 L	0 LM	0 RM	5 R
30	T	67;46	15;23	43;31	61;16
6	B	32;86	56;58	38;29	89;62

Figure 7: Games from Costa-Gomes et al. (2001).