

Stationary Concepts for Experimental 2x2-Games

By REINHARD SELTEN AND THORSTEN CHMURA*

Five stationary concepts for completely mixed 2x2-games are experimentally compared: Nash equilibrium, quantal response equilibrium, action-sampling equilibrium, payoff-sampling equilibrium (Martin J. Osborne and Ariel Rubinstein 1998), and impulse balance equilibrium. Experiments on 12 games, 6 constant sum games, and 6 nonconstant sum games were run with 12 independent subject groups for each constant sum game and 6 independent subject groups for each nonconstant sum game. Each independent subject group consisted of four players 1 and four players 2, interacting anonymously over 200 periods with random matching. The comparison of the five theories shows that the order of performance from best to worst is as follows: impulse balance equilibrium, payoff-sampling equilibrium, action-sampling equilibrium, quantal response equilibrium, Nash equilibrium. (JEL C70, C91)

Experimental evidence suggests that mixed Nash equilibrium is not a very good predictor of behavior. Thus, Ido Erev and Alvin E. Roth (1998, 853) conclude as their first summary observation that “... in some of the games the equilibrium prediction does very badly.” A normal form game is called completely mixed if it has only one equilibrium point in which every pure strategy is used with positive probability. Of special interest are 2x2-games of this kind. They are the simplest games for which mixed equilibrium is the unequivocal game theoretic prediction, if they are played as noncooperative one-shot games.

Mixed equilibrium has several interpretations. One interpretation is that of a rational recommendation for a one-shot game. Another interpretation looks at mixed equilibrium as a result of evolutionary or learning processes in a situation of frequently repeated play with two populations of randomly matched opponents. One may speak of mixed equilibrium as a behavioral stationary concept. Ken Binmore, Joe Swierzbinski, and Chris Proulx (2001) argue in their paper that mixed Nash equilibrium predicts reasonably well for completely mixed, constant sum 2x2-games. However, it is difficult to judge the goodness of fit, if there is no comparison to other stationary concepts.

Economic theory makes extensive use of the concept of mixed equilibrium. One of its attractions is its independence of parameters outside the structure of the game. For the purpose of analyzing theoretical models, it is of great advantage to be able to rely on stationary concepts.

In this paper we will present several alternative stationary concepts for 2x2-games, which can be compared with mixed equilibrium and with each other. For this purpose, we have performed experiments on 12 completely mixed 2x2-games. Six of them are constant sum games and the other six are nonconstant sum games. Each of the constant-sum games was run with 12 independent subject groups and each of the other games with 6 independent subject groups. Each independent subject group consisted of four players 1 and four players 2, interacting in fixed roles

* Selten: Laboratory of Experimental Economics, Adenauerallee 24-42, 53113 Bonn, Germany (e-mail: rselten@uni-bonn.de); Chmura: Shanghai Jiao Tong University, Department of Economics, Fa Hua Zhen Road No. 535, Shanghai 200052, Peoples Republic of China (e-mail: chmura@uni-bonn.de). Financial support by the Deutsche Forschungsgemeinschaft is gratefully acknowledged. We also want to thank Sebastian Goerg for his help. Moreover, we are grateful for the advice given by the editor and three unknown referees.

over 200 periods with random matching. The stationary concepts compared were: Nash equilibrium, quantal response equilibrium, action-sampling equilibrium, payoff-sampling equilibrium, and impulse balance equilibrium.

Quantal response equilibrium (Richard D. McKelvey and Thomas R. Palfrey 1995) assumes that players give quantal best responses to the behavior of the others (see Section IB). In the exponential form of quantal response equilibrium considered here, the probabilities are proportional to an exponential with the expected payoff times a parameter in the exponent.

Action-sampling equilibrium is based on the idea that, in a stationary situation, a player takes a sample of seven observations of the strategies played on the other side, and then optimizes against this sample. If a player has a unique pure best response to her sample, then she plays this strategy. If both strategies are best responses, then each of them is chosen with probability $\frac{1}{2}$. This yields a mixed strategy depending on the probabilities of pure strategies on the other side. Action-sampling equilibrium is a mixed strategy combination consistent with this picture. The name “action-sampling equilibrium” refers to the sampling of the opponent’s actions. The concept has been developed by one of the authors (Selten). As far as we know, it cannot be found in the literature. However the sampling of actions of other players also appears in a paper by Osborne and Rubinstein (1993) in the context of a sampling equilibrium for a large voting game. The sample size is a parameter. Originally the sample size 7 was chosen in view of the famous paper “The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information” by George A. Miller (1956). Later we found that seven actually gives a better fit than other sample sizes.

Payoff-sampling equilibrium (Osborne and Rubinstein 1998) envisions a stationary situation in which a player takes two samples of equal size, one for each of her pure strategies. She then compares the sum of her payoffs in the two samples and plays the strategy with the higher payoff sum. If both payoff sums are equal, then both pure strategies are chosen with probability $\frac{1}{2}$. Payoff-sampling equilibrium is a mixed strategy combination reflecting this picture. Here, too, the sample size is a parameter. The best fitting sample size turns out to be six for each of both samples. The name “payoff-sampling equilibrium” refers to the sampling of own payoffs for each pure strategy.

Impulse balance equilibrium proposed by one of the authors (Selten) is based on learning direction theory (Selten and Joachim Buchta 1999). This learning theory is applicable to the repeated choice of the same parameter in learning situations in which the decision maker receives feedback, not only about the payoff for the choice taken, but also for the payoffs connected to alternative actions. If a higher parameter would have brought a higher payoff, we speak of an upward impulse, and if a lower parameter would have yielded a higher payoff, we speak of a downward impulse. The decision maker is assumed to have a tendency to move in the direction of the impulse.

It is worth pointing out that impulse learning is very different from reinforcement learning. In reinforcement learning, the payoff obtained for a pure strategy played in the preceding period determines the increase of the probability for this strategy. The higher this payoff, the greater is this increase. In impulse learning it is not the payoff in the preceding period that is of crucial importance. It is the difference between what could have been obtained and what has been received, which moves the behavior in the direction of the higher payoff. Moreover, reinforcement learning is entirely based on observed own payoffs, whereas impulse learning requires feedback on the other player’s choice and the knowledge of the player’s own payoff.

In Selten, Klaus Abbink, and Ricarda Cox (2005) impulse balance theory, a semi-quantitative version of learning direction theory, has been proposed. The learning process itself is not modeled, but only the stationary distribution. In the stationary distribution, expected upward impulses are equal to expected downward impulses. As in prospect theory (Daniel Kahnemann and Amos Tversky 1979), losses are counted double in the computation of impulses (formally, this involves the computation of a loss impulse).

Impulse balance equilibrium applies the idea of impulse balance theory to 2×2 -games. The probability of choosing one of two pure strategies, say strategy *A*, is looked upon as the parameter to be adjusted upward or downward. It is assumed that the pure strategy maximin is the reference level determining what is perceived as profit or loss. In impulse balance equilibrium, expected upward and downward impulses are equal for each of both players simultaneously.

Following a suggestion of one of the authors (Selten), impulse balance equilibrium has been successfully applied to special 2×2 - and $2 \times 2 \times 2$ -games in a paper by Judith Avrahami, Werner Güth, and Yaakov Kareev (2005).

Remarks: Two of the stationary concepts compared in this paper, Nash equilibrium and impulse balance equilibrium, are parameter free. Action-sampling equilibrium involves one parameter, namely, the number seven which, however, has been chosen in view of admittedly quite weak theoretical considerations confirmed by pilot experiments not included in the main sample of this paper. A similar theoretical reasoning would suggest seven as the sample size for payoff-sampling equilibrium. However, there six yields the best fit to the data. Quantal response equilibrium involves one parameter, namely, the constant multiplier of expected payoffs in the exponent. This parameter has to be adjusted to the data. There are no theoretical considerations, not even very weak ones, which could be used in order to determine this parameter in any other way.

Quantal response equilibrium modifies Nash equilibrium by introducing noise into the optimization process. Thereby, the best response notion is replaced by a notion of quantal response. The two sampling equilibria, action-sampling equilibrium and payoff-sampling equilibrium, also involve noise produced by sampling error. In contrast to quantal response equilibrium, however, this noise is endogenous and is completely determined by the sample size and the payoffs of the game.

Quantal response equilibrium is not connected to any theory that relates the noise parameter to the structure of the game. One could, of course, fit the parameter for every individual game separately. However, this does not yield a method for predicting a unique stationary mixed strategy combination for every completely mixed 2×2 -game. In order to make the concept of quantal response equilibrium comparable to other theories involving at most one parameter, one has to look at the parameter of quantal response equilibrium as an unknown behavioral constant which is the same for all games. Accordingly, we determine the value of the parameter that best fits all our data, and base our comparison on this.

The five concepts can be thought of as stationary states of dynamic learning models. Learning models differ with respect to their requirements on prior knowledge of the game and on feedback after each period. Nash equilibrium is stationary with respect to reinforcement learning models like the ones used by Erev and Roth (1998). These models require feedback on own payoffs but not more. A player does not even have to know his or her own payoff matrix. The same knowledge and feedback requirements are sufficient for learning models with quantal response equilibrium as stationary state. The expected payoffs appearing in the formulas for quantal response equilibrium can be estimated as average past payoffs. Simple learning models yielding payoff-sampling equilibrium as stationary state immediately suggest themselves. It is clear that here, too, only feedback of a player's own period payoff is necessary.

The other two concepts seem to be more demanding with respect to learning models yielding them as stationary states. As far as we can see, one needs knowledge of one's own payoff matrix, as

well as feedback on the other player's choice in these two cases. Clearly, a player must know his or her own payoff matrix for optimizing against a sample of the other player's choices. The same kind of knowledge and feedback is necessary for perceiving impulses in learning direction theory.

The development of stationary concepts that fit experimental data is very important for behavioral theory. With the help of such concepts, theoretically interesting situations can be mathematically explored as, for example, a voting situation in a paper by Osborne and Rubinstein (2003).

Learning models could also be applied to theoretically interesting situations. However, the construction of learning models usually involves many details which may influence the outcome of computer simulations. This makes it difficult to work with learning models rather than stationary concepts. Moreover, in complex situations, one may need a huge number of computer simulations in order to answer questions of comparative statics, which can be attacked mathematically on the basis of stationary concepts.

In completely mixed 2x2-games, Nash equilibrium and impulse balance equilibrium can be described by explicit formulas, and therefore are easy to use in theoretical investigations. This is not true, however, for quantal response equilibrium, action-sampling equilibrium, or payoff-sampling equilibrium. The latter concepts can be computed numerically only with the help of a computer. Nevertheless it is maybe sometimes possible to investigate their comparative static properties by mathematical operations like implicit differentiation applied to the defining equations. A similar approach to the results of learning models seems to be almost hopeless.

In this paper, all five stationary concepts will be defined only for completely mixed 2x2-games. In the literature, Nash equilibrium, quantal response equilibrium, and payoff-sampling equilibrium are defined for normal form games in general. It is also clear how the concept of action-sampling equilibrium can be generalized to all normal form games. Admittedly, this is less clear for impulse balance equilibrium as far as normal forms with more than 2 strategies for some players are concerned. Here, different generalizations are possible. The basic principle would be that for each strategy of a player, expected incoming impulses should be equal to expected outgoing impulses unless there are no outgoing impulses, as in pure Nash equilibrium. In Appendix F (part of the online Appendix, available at <http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.3.938>) a sketch of a generalization of impulse balance equilibrium to general n -person games in normal form is presented.

The comparison of stationary concepts can also guide the search for adequate learning rules. In the past, many authors, like Selten (1990) and Sergiu Hart and Andreu Mas-Colell (2000), felt that a reasonable learning model should converge to Nash equilibrium or correlated equilibrium under favorable assumptions. If, however, other stationary concepts better fit experimental data, one may want to look at learning processes converging to them.

As we shall see, over all 200 periods and all 108 independent subject groups, the comparison yields the following order with respect to the goodness of fit from best to worst: impulse balance equilibrium, payoff-sampling equilibrium, action-sampling equilibrium, quantal response equilibrium, Nash equilibrium. However, the difference between impulse balance equilibrium and payoff-sampling equilibrium is not statistically significant (see Section IIIH).

In Section I we shall present a more detailed description of the five concepts. Section II will explain the experimental setup, and Section III will describe the results. Section IV concludes with a summary and discussion.

I. The Five Stationary Concepts

All the experimental 2x2-games in this paper have the structure shown by Figure 1. The arrows around the matrix show the direction of best replies. The parameters a_L , a_R , b_U , and b_D are assumed to be nonnegative. Games with negative payoffs probably would require special

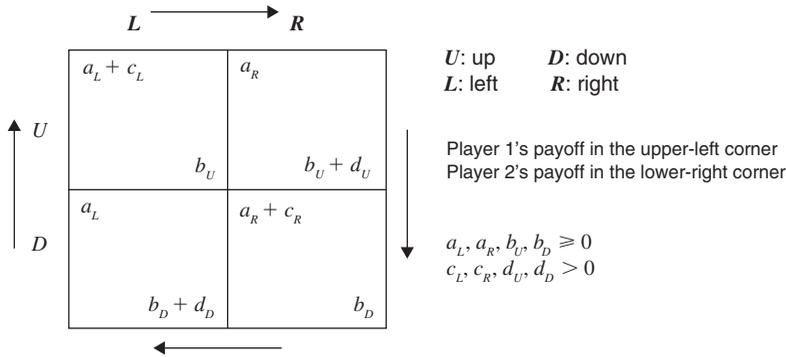


FIGURE 1. STRUCTURE OF THE EXPERIMENTAL 2X2-GAMES

behavioral considerations, which we want to avoid in this paper. The parameters c_L and c_R are player 1's payoff differences in favor of U and D , respectively. Similarly, d_U and d_D are payoff differences of player 2 for R and L , respectively. All these payoff differences are assumed to be positive. It is clear that a game with this structure is completely mixed, in the sense that it has a uniquely determined, completely mixed Nash equilibrium.

In a completely mixed 2x2-game, the arrows may also have the opposite orientation. However, we can restrict our attention to the structure shown by Figure 1 without any loss of generality. The case of counterclockwise arrows can be transformed to the one shown above by an interchange of the two rows.

A. Equilibrium Conditions and Their Graphical Representation

Let $p = (p_U, p_D)$ and $q = (q_L, q_R)$ be the mixed strategies of player 1 and player 2, respectively. Here, p_U and p_D are player 1's choice probabilities for U and D , and q_L and q_R are player 2's choice probabilities for strategy L and R . The space of mixed strategies for a game with a structure of Figure 1 can be described by the (p_U, q_L) -diagram, which shows the interval $0 \leq p_U \leq 1$ horizontally and the interval $0 \leq q_L \leq 1$ vertically. Every point (p_U, q_L) in this square represents a strategy combination.

Each of the five concepts involves two equilibrium conditions. The first one describes equilibrium adjustment of player 1 for any given mixed strategy of player 2. In the same way, the second condition expresses equilibrium adjustment of player 2 to any given mixed strategy of player 1. These two equilibrium conditions can be represented by curves in the (p_U, q_L) -diagram. We call the graph of the first equilibrium condition the curve for p_U and the graph for the second one the curve for q_L . The intersection of both curves is the stationary equilibrium specified by the concerning concept.

Figure 2 shows the curves for p_U and q_L arising in the example of our experimental game 1 (see Figure 5 in IIB). With the exception of the case of Nash equilibrium, the curves for p_U are monotonically increasing and the curves for q_L are monotonically decreasing. In all five parts of Figure 2, both curves intersect at the relevant stationary equilibrium of our experimental game 1.

We now briefly discuss the two curves in the case of the Nash equilibrium. Let p_U^N and p_L^N be the Nash equilibrium probabilities for U and L , respectively. Let us look at p_U on the curve for p_U as q_L moves from zero to one. In the first vertical piece of the curve with $0 \leq q_L \leq q_L^N$, the probability p_U remains constant at $p_U = 0$. Then it moves on a horizontal piece at q_L^N from zero to one. The curve ends with a vertical piece with $q_L^N \leq q_L \leq 1$, at which p_U stays at $p_U = 1$. Similarly, on the curve for q_L , the probability q_L stays at $q_L = 1$ in a horizontal piece with $0 \leq q_U \leq p_U^N$, then

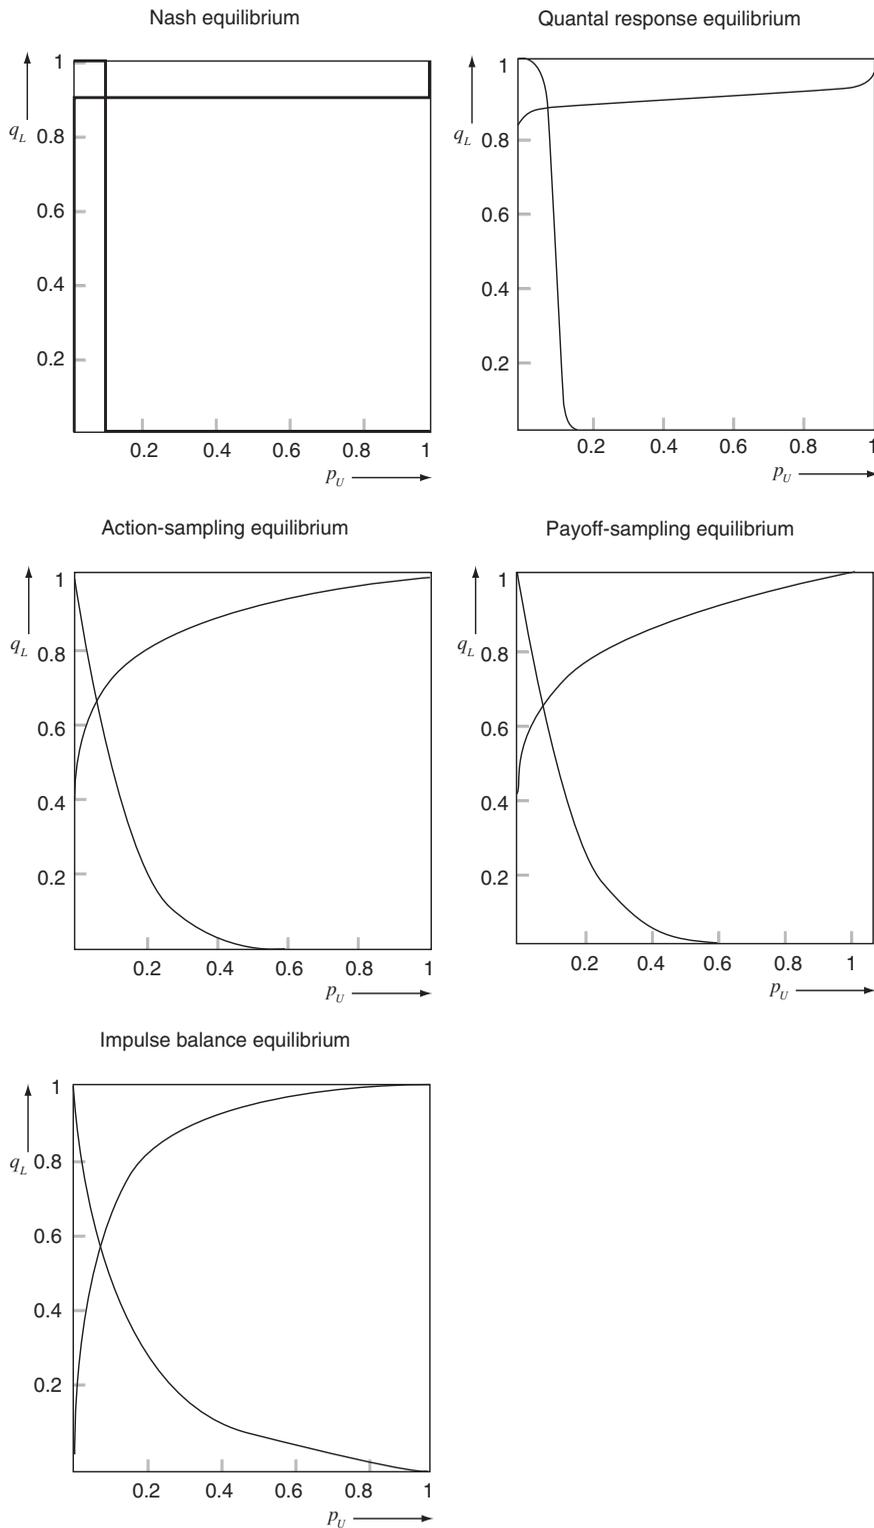


FIGURE 2. THE CURVES FOR p_U AND q_L ARISING IN THE EXAMPLE OF GAME 1 FOR EACH OF THE FIVE CONCEPTS

decreases from one to zero on a vertical piece with $q_L = q_L^N$, and finally comes to a horizontal piece with $q_L^N \leq q_L \leq 1$ and $p_U = 0$. In this sense, one may say that p_U is increasing or constant along the curve for p_U , and q_L is decreasing or constant along the curve for q_L .

In the case of the other four concepts, the curves for p_U and q_L are continuously differentiable. For each of these concepts, equations for the two curves will be given in the Section II B, C, D, and E. In these cases, the value of p_U at q_L on the curve for p_U is denoted by $p_U(q_L)$. Similarly, the notation $q_L(p_U)$ is used for the value of q_L at p_U on the curve for q_L .

The curves for the concepts different from Nash equilibrium reveal a considerable sensitivity with respect to the strategy of the other player. Suppose, for example, player 2 plays her Nash equilibrium strategy q_L^N and player 1 chooses the strategy $p_U(q_L^N)$. The value of $p_U(q_L^N)$ for quantal response equilibrium, action-sampling equilibrium, payoff-sampling equilibrium, and impulse equilibrium is 0.29, 0.52, 0.56, and 0.33, respectively, whereas p_U^N is equal to 0.09. It can be seen that in all four cases there is a considerable difference between $p_U(q_L^N)$ and p_U^N .

A look at Figure 2 suggests a distinction of two groups of the pictures shown there. The first group consists of the two diagrams in the first row and the second group is formed by the remaining three pictures. The curves for quantal response equilibrium are near to those of Nash equilibrium. In this respect, there is a close similarity within the first group. The diagrams within the second group also look very similar to each other, but there is a marked difference between the two groups.

As we shall see later, the concepts giving rise to the second group of pictures clearly outperform those connected to the first group. These three concepts yield predictions near to each other and much nearer to the observed relative frequencies.

In online Appendix D, it will be shown for each of the five stationary concepts that the curves for p_U and q_L always have a unique intersection. Therefore, the stationary equilibrium exists and is uniquely determined in all five cases.

In completely mixed games, the Nash equilibrium strategy of a player is independent of his own payoff. As one would intuitively expect, experimental findings suggest that an increase of a player's payoff in one of the four fields with all other payoffs of both players kept constant tends to increase the probability of this player's strategy used in this field. In online Appendix E it will be shown that, at equilibrium, such payoff changes always increase this probability for quantal response equilibrium and for impulse balance equilibrium and, in the case of action-sampling equilibrium and payoff-sampling equilibrium, this probability is never decreased, but increased if the payoff change is big enough. The two sampling equilibria depend discontinuously on payoffs.

B. Nash Equilibrium

In the case of Nash equilibrium, the curves for p_U and q_L are the graphs of the best reply correspondences for the two players (see Figure 2). The choice probabilities are as follows:

$$(1) \quad p_U = \frac{d_D}{d_U + d_D}, \quad p_D = \frac{d_U}{d_U + d_D}, \quad q_L = \frac{c_R}{c_L + c_R}, \quad q_R = \frac{c_L}{c_L + c_R}.$$

The choice probabilities of a player in Nash equilibrium are independent of his own payoff. They are entirely determined by the payoff differences of the other player. This is a well-known counterintuitive property of Nash equilibrium.

C. Quantal Response Equilibrium

It is assumed that players choose a "quantal best response" to the strategies of the other player. They make mistakes, taking the mistakes of the other player into account.

Let $E_U(q)$ and $E_D(q)$ be player 1's expected payoff for U and D , respectively, against a strategy q of player 2. Similarly, $E_L(p)$ and $E_R(p)$ are player 2's expected payoffs for L and R , respectively, against a strategy p of player 1.

In quantal response equilibrium, the curves for p_U and q_L are as follows:

$$(2) \quad p_U = \frac{e^{\lambda E_U(q)}}{e^{\lambda E_U(q)} + e^{\lambda E_D(q)}}, \quad q_L = \frac{e^{\lambda E_L(p)}}{e^{\lambda E_L(p)} + e^{\lambda E_R(p)}}.$$

These equations yield a simultaneous equation system, which determines the choice probabilities as functions of λ . For our data, $\lambda = 8.84$ is the best fitting overall estimate. This value of λ minimizes the sum of mean squared distances from the actually observed relative choice frequencies for the 12 experimental games. This measure of predictive success will be explained in Section IIIB.

The best response structure of a two-person game is a pair of mappings (α, β) . The mapping α maps the strategies q of player 2 to player 1's set $\alpha(q)$ of pure best responses to q , and the mapping β maps the mixed strategies p of player 1 to the set $\beta(p)$ of player 2's pure best responses to p . Nash equilibrium depends only on the best response structure of the game. However, quantal response equilibria with the same parameter λ can be different for two games with the same best response structure. If all payoffs of a 2x2-game are multiplied by the same positive factor x , the best response structure remains unchanged, but quantal response equilibrium for a fixed parameter λ does change. The multiplication of all payoffs by x has the same effect as not changing payoffs and replacing λ by $\lambda' = \lambda x$.

Suppose that the payoffs are changed by adding a constant r to all payoffs of player 1 in row R of Figure 1 and leaving everything else unchanged. Let $E'_U(q)$ and $E'_D(q)$ be the new payoffs for U and D in the new game obtained in this way. We have

$$(3) \quad E'_U(q) = E_U(q) + q_R r, \quad E'_D(q) = E_D(q) + q_R r.$$

This means that the equation for p_U in the new game can be simplified by dividing numerator and denominator by the common factor $e^{q_R r}$. Therefore, the equations for p_U and p_D do not really change in the transition to the new game. The same argument can be applied to the case that a constant is added to player 1's payoff in the column L or player 2's payoff in one of the two rows. We can conclude that such additive changes do not have any effect on the quantal response equilibrium, even if it does not depend on the best response structure alone.

D. Action-Sampling Equilibrium

In the stationary state described by p_U, p_D, q_L , and q_R , player 1 takes a sample of n choices L or R and optimizes against this sample. Player 2 behaves analogously. This concept describes a stationary state of two large populations of players 1 and 2. Every member takes a sample of n past decisions of players on the other side and optimizes against it. More precisely, he chooses his best response if this is uniquely determined and plays his mixed strategy $(\frac{1}{2}, \frac{1}{2})$ if both pure strategies are best responses. The action-sampling equilibrium is a stationary state of this system. Here, too, p_U, p_D, p_L , and p_R are stationary probabilities of U, D, R , and L . Consider two specific players 1 and 2 in both populations. Let k be the number of L 's in player 1's sample and let m be the number of D 's in player 2's sample. Then, players 1 and 2 will play as follows:

- Player 1 plays $U, D, (\frac{1}{2}, \frac{1}{2})$ for $kc_L > (n - k)c_R, kc_L < (n - k)c_R, kc_L = (n - k)c_R b$, respectively;

- Player 2 plays $L, R, (\frac{1}{2}, \frac{1}{2})$ for $md_D > (n - m)d_U, md_D < (n - m)d_U, md_D = (n - m)d_U$, respectively.

Instead of $kc_L > (n - k)c_R$, we also can write

$$(4) \quad \frac{k}{n} > \frac{c_R}{c_L + c_R}.$$

Let $\alpha_U(k)$ be the probability of player 1 choosing U for k and $\alpha_L(m)$ be the probability of player 2 choosing L for m .

It can be seen immediately that we have

$$(5) \quad \alpha_U(k) = \begin{cases} 1 & \text{for } \frac{k}{n} > \frac{c_R}{c_L + c_R} \\ \frac{1}{2} & \text{for } \frac{k}{n} = \frac{c_R}{c_L + c_R} \\ 0 & \text{else} \end{cases}, \quad \alpha_L(m) = \begin{cases} 1 & \text{for } \frac{m}{n} > \frac{d_U}{d_U + d_D} \\ \frac{1}{2} & \text{for } \frac{m}{n} = \frac{d_U}{d_U + d_D} \\ 0 & \text{else} \end{cases}.$$

L is played with the probability q_L . Accordingly, the number k of L 's in player 1's sample is binomially distributed. An analogous statement holds for the number of D 's in player 2's sample. One obtains the following equations for p_U and q_L .

$$(6) \quad p_U = \sum_{k=0}^n \binom{n}{k} q_L^k (1 - q_L)^{n-k} \alpha_U(k), \quad q_L = \sum_{m=0}^n \binom{n}{m} (1 - p_U)^m p_U^{n-m} \alpha_L(m).$$

These equations describe the curves for p_U and q_L explained in Section IA.

Remarks: The functions $\alpha_U(k)$ and $\alpha_L(m)$ depend only on the payoff differences c_L, c_R, d_U , and d_D . Therefore, the concept of action-sampling equilibrium depends only on the best response structure.

The curves for p_U and q_L are differentiable with respect to q_L and p_U , respectively, for given payoff differences c_L, c_R, d_U , and d_D ; however, the two curves do not depend continuously on these payoff differences. If, for example, $c_R/(c_L + c_R)$ is equal to $\frac{1}{2}$, a small change of either c_L or c_R results in a jump of $\alpha_U(k)$.

The concept of action-sampling equilibrium can easily be extended to general normal form games. In a stationary situation, a player takes a sample of seven observations of combinations of pure strategies for the other players and then optimizes against this sample. In the case of several best responses, each of them is chosen with equal probability.

E. Payoff-Sampling Equilibrium

The basic idea of payoff-sampling equilibrium has been explained in the introduction. Osborne and Rubinstein (1998) did not specify the probabilities of both strategies in the case that the payoff sums for the two samples are equal. In order to obtain a unique prediction, we added the rule that, in this case, each pure strategy is chosen with probability $\frac{1}{2}$.

As before, p_U, p_D, q_L , and q_R denote the stationary probability for the corresponding pure strategies.

Let n be the sample size and k_U and k_D be the number of L 's in a player 1's sample for U and D , respectively. Similarly, let m_L and m_R be the number of U 's in a player 2's sample for L and R , respectively.

Player 1's sums of payoffs H_U and H_D in the samples for U and D , respectively, are as follows:

$$(7) \quad H_U = k_U(a_L + c_L) + (n - k_U)a_R, \quad H_D = k_D a_L + (n - k_D)(a_R + c_R).$$

In the same way, player 2's sum of payoffs in the samples for L and R are given by

$$(8) \quad H_L = m_L b_U + (n - m_L)(b_U + d_D), \quad H_D = m_R(b_U + d_U) + (n - m_R)b_D$$

Player 1's probability $\beta_U(k_U, k_D)$ of playing U if k_U and k_D are the numbers of L 's in his sample, as well as the probability $\gamma(m_L, m_R)$ of player 2 playing L if she observes the numbers m_L and m_R of U 's in her samples for L and R , are as follows:

$$(9) \quad \beta(k_U, k_D) = \begin{cases} 1 & \text{for } H_U > H_D \\ \frac{1}{2} & \text{for } H_U = H_D, \\ 0 & \text{else} \end{cases}, \quad \gamma(m_L, m_R) = \begin{cases} 1 & \text{for } H_L > H_R \\ \frac{1}{2} & \text{for } H_L = H_R. \\ 0 & \text{else} \end{cases}$$

Since k_U and k_D , as well as m_L and m_R , are binomially distributed, we have

$$(10) \quad p_U = \sum_{k_U=0}^n \sum_{k_D=0}^n \binom{n}{k_U} \binom{n}{k_D} q_L^{k_U+k_D} (1 - q_L)^{2n-k_U-k_D} \beta(k_U, k_D).$$

$$(11) \quad q_L = \sum_{m_L=0}^n \sum_{m_R=0}^n \binom{n}{m_L} \binom{n}{m_R} (1 - p_U)^{m_L+m_R} p_U^{2n-m_L-m_R} \gamma(m_L, m_R).$$

The curves for p_U and q_L in the case of payoff-sampling equilibrium are represented by these two equations.

Remarks: The operation of adding a constant to player 1's payoffs in the column for R may change $\beta_U(k_U, k_D)$ and, therefore, the first of the two equations. Similarly, adding a constant to player 2's payoffs in the row for U may change the second equation. For this reason, payoff-sampling equilibrium is not invariant with respect to these operations.

As in the case of action-sampling equilibrium, the curves for p_U and q_L are differentiable with respect to q_L and p_U , respectively, but not continuous with respect to a small change of one payoff in the payoff-matrix for the concerning player.

F. Impulse Balance Equilibrium

As was explained in the introduction, impulse balance theory is not applied to the original game, but to a transformed game, in which losses with respect to a natural aspiration level get twice the weight as gains above this level.

The natural aspiration level for a player is his pure strategy maximin value or, in other words, the maximum of the lowest payoff he may obtain for using one of his pure strategies. Define:

$$(12) \quad s_1 = \max[\min(a_L + c_L, a_R), \min(a_L, a_R + c_R)].$$

$$(13) \quad s_2 = \max[\min(b_U, b_D + d_D), \min(b_U + d_U, b_D)].$$

From now on, we shall refer to s_1 and s_2 as the pure strategy maximin payoffs or shortly the security levels of players 1 and 2, respectively.

In the following, it will be argued that the security level of a player is her second lowest payoff. It may happen that the lowest payoff is obtained at more than one of the four fields. In this case, there is no difference between the second lowest payoff and the lowest payoff. The words “second lowest payoff” will always be understood this way.

In a completely mixed 2x2-game, no pure strategy can dominate another one (see Figure 1). Therefore, the lowest and the second lowest payoff of player 1 cannot appear in the same row. An analogous statement holds for player 2.

The second lowest payoff is always at least obtained if it is the lowest one. Otherwise, the lowest payoff can be avoided by not choosing the pure strategy which may yield it. Thereby, the second lowest payoff is secured. It is also clear that one cannot secure more than that by the use of a pure strategy.

The security level can be enforced, no matter what the other player does. Therefore, it is natural to look at a lower payoff as a failure, and its difference to the security level as a loss. It makes no sense to be satisfied with less than one could have had for sure. Loss aversion is a well-known behavioral concept used, for example, in prospect theory (Kahnemann and Tversky 1979). In the case of a payoff below the security level, there are two reasons for thinking that one should have chosen the other strategy. The first is that the other strategy would have yielded a higher payoff. The second is that the loss should be avoided. The loss counts as a part of the foregone payoff and, in addition, counts once more by its quality of being a loss rather than merely a foregone gain.

An earlier formulation of impulse balance theory concerned an auction situation in which losses could occur only in connection with bids appearing to be too high ex post (Selten, Abbink, and Cox 2005). Therefore, in the case of a loss, the decision maker experienced a downward impulse and a loss impulse. In 2x2-games, losses may occur for choices of one strategy or the other, depending on the structure of the game. Thus, in game 1 (see Figure 5, p. 951), player 1 at (U, R) experiences a loss of nine and a foregone payoff of ten. Therefore, a loss impulse of nine is added to the ordinary impulse of ten from U to D at (U, R) . At (D, L) player 1 receives only an ordinary impulse of one from D to U .

As we shall see, the combination of ordinary impulses and loss impulses is automatically taken care of if impulses from one pure strategy to another are computed in a transformed game in which losses receive double the weight of gains. We construct this transformed game by leaving player i 's payoffs below and at s_i unchanged, and by reducing the surplus over s_i of higher payoffs by the factor $\frac{1}{2}$. Figure 3 shows the *impulse balance transformation* for the example of experimental game 3 (see Figure 5).

The payoff differences in the transformed game corresponding to c_L, c_R, d_U, d_D are denoted by $c_L^*, c_R^*, d_U^*, d_D^*$. If, after a play, player i could have obtained a higher payoff by the choice of his other strategy, he receives an “impulse” in the direction of his other strategy. The size of this impulse is the foregone payoff in the transformed game. If, for example, player 1 chooses U and the other player chooses R , then player 1 receives an impulse of $c_R^* = 8.5$ in the direction of D . A player receives no impulse if the payoff for the strategy he did not choose was lower than the one he obtained. Figure 4 shows the impulses in the direction of the strategy not chosen, similar to a payoff table.

It can now be seen without difficulty that impulses in the transformed game automatically combine ordinary impulses and loss impulses in the original game. In the case of a payoff below s_i , the loss part of an impulse is fully counted and a possible foregone-gain part is reduced by the factor $\frac{1}{2}$, just like an impulse in the case of a payoff above the security level. Half of the fully counted loss corresponds to the loss impulse.

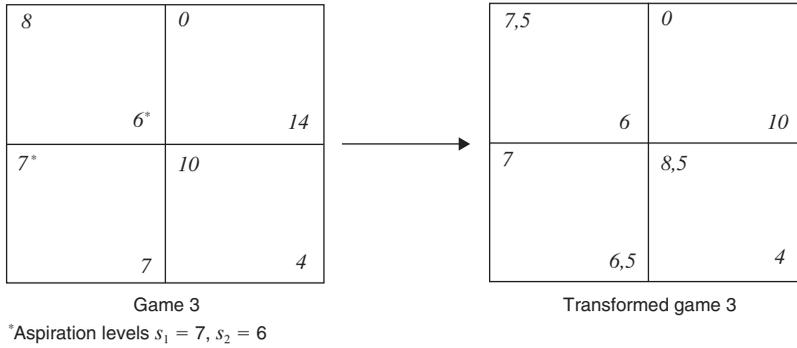


FIGURE 3. IMPULSE BALANCE TRANSFORMATION FOR THE EXAMPLE OF EXPERIMENTAL GAME 3

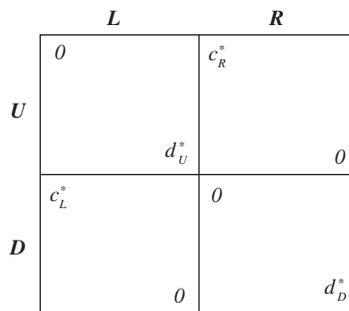


FIGURE 4. IMPULSE IN THE DIRECTION OF THE STRATEGY NOT CHOSEN

Impulse balance equilibrium requires that player 1’s expected impulse from *U* to *D* is equal to his expected impulse from *D* to *U*. Similarly, player 2’s expected impulse from *L* to *R* must be equal to her expected impulse from *R* to *L*. This yields the following two *impulse balance equations*:

$$(14) \quad p_U q_R c_R^* = p_D q_L c_L^*, \quad p_U q_L d_U^* = p_D q_R d_D^*.$$

The left-hand side of the first impulse balance equation is player 1’s expected impulse from *U* to *D*, and the right-hand side is player 1’s expected impulse from *D* to *U*. If the left-hand side is greater than the right-hand side, then player 1 receives stronger impulses from *R* to *D*, and this will decrease q_R and increase q_L . This creates a tendency in the direction of impulse balance. An analogous interpretation can be given to the second impulse balance equation. Of course this is only a heuristic argument. In this paper, we do not want to explore the dynamics of impulse balance equilibrium.

The impulse balance equations yield the following equations of the curves for p_U and q_U :

$$(15) \quad p_U = \frac{q_L c_L^*}{q_L c_L^* + (1 - q_L) c_R^*}, \quad q_L = \frac{(1 - p_U) d_D^*}{p_U d_U^* + (1 - p_U) d_D^*}.$$

In Section E4 of online Appendix E, explicit formulas will be derived for the coordinates of the intersection (p_U, q_L) of the two curves. Define $c = c_L^*/c_R^*$ and $d = d_U^*/d_D^*$: it will be shown in E4 that at the intersection we have $p_U = \sqrt{c}/(\sqrt{c} + \sqrt{d})$ and $q_L = 1/(1 + \sqrt{cd})$.

In online Appendix F a possibility of generalizing impulse balance equilibrium to n -person normal form games will be briefly sketched. Even if for the substance of this paper no such generalization is needed, it is maybe of interest to see in which way it could be achieved.

II. Experimental Design

A. Procedure

The experimental data were obtained in 54 sessions with 16 subjects each, 864 altogether. The subjects were students of the University of Bonn, mainly majoring in economics or law. The experiments were run in the Bonn Experimental Economics Laboratory. The computer program was based on the toolbox RatImage developed by Abbink and Abdolkarim Sadrieh (1995). Only one game was played in each session.

At the beginning of a session, oral and written instructions were given to the subjects. The written instructions (in German) are shown in online Appendix B. The subjects were informed about the game matrix, including the payoffs of both players. They were told that they would interact with randomly changing opponents and always be in the same player role over 200 periods. Actually, in each session there were two independent subject groups with four participants in the role of player 1 and four participants in the role of player 2. The players played against randomly chosen opponents but only within their independent group. They were not informed about the fact that there are two groups. We did not lie to them, but did convey the impression that they would interact directly or indirectly with 15 other players.

After the instruction, participants sat in separate cubicles and made their decisions by mouse click. The decisions in a play were made without any information about the choices of the other players. After each of the 200 plays, they received feedback about the other players' choice and payoff, period number, and their cumulative payoff. No limit was imposed on the decision time. The subjects were not permitted to take notes of any kind about their playing experience. They were also not permitted to talk to each other during the experiment and they had no opportunity to see the screens of other participants. After each experiment, participants had to fill in a questionnaire. (Because no use of the questionnaire data is made in this paper, the questionnaire is not shown here.)

Each participant received 5€ and, in addition, a monetary payoff proportional to his or her game payoff accumulated over the 200 periods. The exchange rate was 1.6€-Cent per payoff point. An experimental session took 1.5 to 2 hours and the average earning of a subject was about 24€ including the fee for showing up.

In some sessions, a digit span test (E. A. Davis 1931; S. C. Della Sala et al. 1999) preceded the game playing. This test is designed to measure the short-term memory size. As we make no use of the data collected by this test in this paper, the details of the digit span test will not be explained here.

B. Experimental Games

Figure 5 shows the 12 games used in our experiment. The constant sum games are shown on the left side of Figure 5 and the nonconstant sum games on the right side. The nonconstant sum game right next to a constant sum game in the Figure 5 has the same best response structure. We say that the two games form a pair. The nonconstant sum game in a pair is derived from the constant sum game in the pair by adding the same constant to player 1's payoff in the column for R and player 2's payoff in the row for U . It is clear that this does not change the best response structure.

Nash equilibrium and action-sampling equilibrium depend only on the best response structure, and therefore yield the same predictions for both games in a pair. In Section IB, we explained that adding a constant to all payoffs of player 1 in a specific column or to all payoffs of player 2

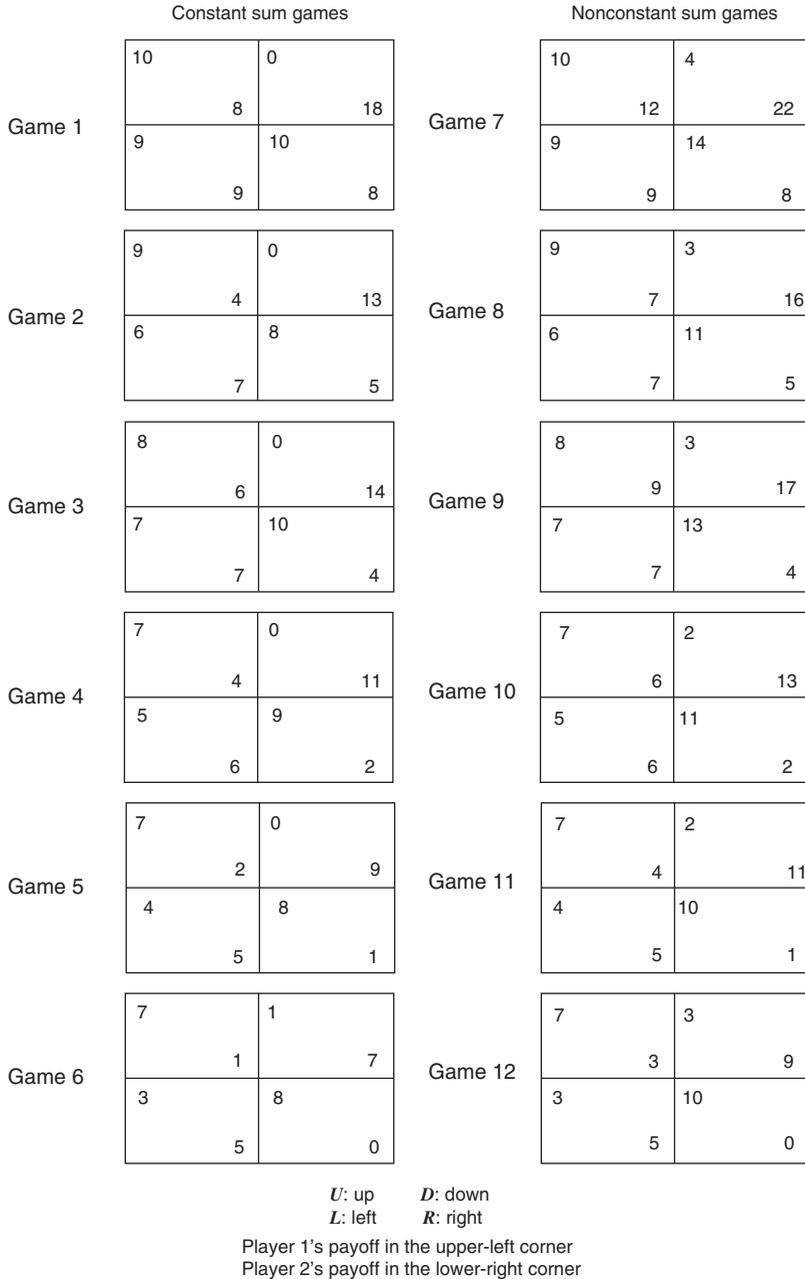


FIGURE 5. EXPERIMENTALLY INVESTIGATED GAMES

in a row does not change the quantal response equilibrium, even if this concept does not depend only on the best response structure. Therefore, quantal response equilibrium also yields the same prediction for the two games in a pair.

The games in a pair also have the same action-sampling equilibrium. A best response to a sample of pure strategies of the other player in one of the two games is also a best response to

this sample in the other game. This is an immediate consequence of the fact that both games have the same best response structure.

In view of the remark at the end of Section ID one cannot expect that payoff-sampling equilibrium generates the same prediction for the games in a pair. In fact, these predictions are different for all six pairs.

The determination of impulse balance equilibrium involves a transition from the original game to the transformed game. The pure strategy maximin payoff, which serves as a reference point for gains and losses, may be different for the two games of the pair, and even if this is not the case, the best response structures will usually be different. In fact, in all six cases, the impulse balance equilibria are different for the two games in a pair.

In the selection of the experimental games, we have been guided by several considerations explained in the following. Two pilot experiments were run with the games shown in Figure 6. Game A is similar to the game played by Jack Ochs (1995) and also by Jakob K. Goeree, Charles A. Holt, and Palfrey (2003). In the questionnaires, the subjects who had played game A often reported attempts to cooperate.

Even if these attempts failed, they may have had an influence on the observed relative frequencies. Therefore, we decided to explore constant sum games extensively. Constant sum games offer no cooperation opportunities. We wanted to contrast them with similar nonconstant sum games offering some scope for cooperation.

The concepts of action-sampling equilibrium and impulse balance equilibrium have been developed on the basis of the pilot experiments with games A and B. Therefore, the experimental results obtained with these games are not included in the comparison of the five theories.

The selection of the constant sum games was guided by the idea that, on the one hand, a reasonably wide distribution over the parameter space should be achieved, and, on the other hand, the number of games should be small enough to permit a sufficiently large number of independent subject groups in every case.

The games explored here have eight payoffs, but the best response structure is characterized by two parameters. The Nash equilibrium choice probabilities p_U^N and q_L^N will serve as these two parameters in the following figure. Figure 7 show the six Nash equilibria for the experimental games.

In all six cases, p_U^N is between 0 and 0.5 and q_L^N is between 0.5 and 1. Therefore only this part of the parameter space is shown in Figure 7. The best reply structure remains essentially unchanged if the rows or columns, or the role of both players, are exchanged. Such transformations yield all the points in Figure 7.

It can be seen that the six games, together with their automorphic transformations, are widely distributed over the parameter space. However, we intentionally underrepresented cases in which one of the equilibrium choice probabilities is near to 0.5. In our sample of six, only game 6 has this property. In the middle of the parameter space, where both parameters are 0.5, every reasonable theory predicts equal probabilities for all strategies. The greater the distance from the midpoint, the more the stationary concepts compared in this paper differ with respect to their predictions.

Since constant sum games are more basic, we have run experiments with 12 independent subject groups for each of the 6 constant sum games but only 6 independent subject groups for each of the nonconstant sum games.

III. Experimental Results

A. Predicted and Observed Relative Frequencies

We begin our descriptions of the results obtained by a number of figures showing the predictions of the five stationary concepts, together with the observed overall relative frequencies, for each of the experimental games. The numerical values are shown in Table 1.

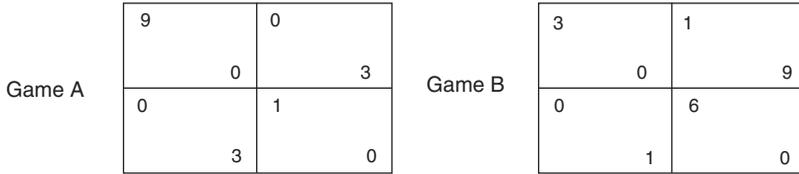


FIGURE 6. STRUCTURE OF THE PILOT EXPERIMENTS

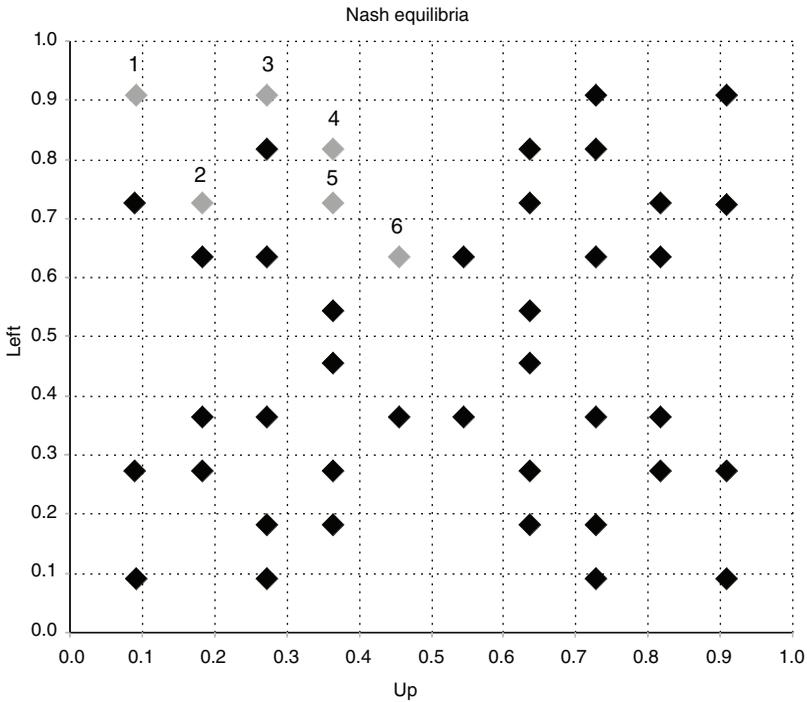


FIGURE 7. PERMUTATIONS OF ROWS, COLUMNS, OR PLAYER ROLES TRANSFORM THE 6 EXPERIMENTAL GAMES INTO 44 GAMES WITH THE NASH EQUILIBRIA SHOWN IN THE FIGURE

In the first three columns of Table 1 the theoretical values of the upper half are repeated in the lower half. This is due to the fact that Nash equilibrium and action-sampling equilibrium depend only on the best response structure (see the remark at the end of Section ID and the property of quantal response equilibrium explained at the end of Section IC).

In Figures 8 and 14 in the online Appendix A, we show cutouts of the whole parameter space with predictions and observed averages for all 12 games. Apart from the fact that the Nash equilibrium of game 2 is nearer to (0.5, 0.5) than that of game 3, the games 1–6 are farther from the middle of the parameter space the lower their order in the numbering. One can see that the discrimination between the concepts tends to be worse for games nearer to the middle of the parameter space. For games 1 to 5, Nash equilibrium and quantal response equilibrium are farther from the observed averages than the other three concepts, but for game 6, all concepts are quite near to the observed average. Since game 6 is near to the middle of the parameter space, random fluctuations seem to play a greater role for this game.

The predictions of impulse balance equilibrium, payoff-sampling equilibrium, and action-sampling equilibrium tend to be near to each other. Therefore random fluctuations make the

TABLE 1—FIVE STATIONARY CONCEPTS TOGETHER WITH THE OBSERVED RELATIVE FREQUENCIES FOR EACH OF THE EXPERIMENTAL GAMES

		Nash equilibrium	Quantal response equilibrium	Action- sampling equilibrium	Payoff- sampling equilibrium	Impulse balance equilibrium	Observed average of 12 observations
Game 1	U	0.091	0.070	0.057	0.071	0.088	0.079
	L	0.909	0.882	0.664	0.645	0.580	0.690
Game 2	U	0.182	0.172	0.185	0.184	0.172	0.217
	L	0.727	0.711	0.619	0.569	0.491	0.527
Game 3	U	0.273	0.250	0.137	0.152	0.161	0.198
	L	0.909	0.898	0.753	0.773	0.765	0.793
Game 4	U	0.364	0.348	0.286	0.285	0.259	0.286
	L	0.818	0.812	0.679	0.726	0.710	0.736
Game 5	U	0.364	0.354	0.286	0.307	0.297	0.327
	L	0.727	0.721	0.679	0.654	0.628	0.664
Game 6	U	0.455	0.449	0.448	0.427	0.400	0.445
	L	0.636	0.634	0.613	0.597	0.600	0.596
		Nash equilibrium	Quantal response equilibrium	Action- sampling equilibrium	Payoff- sampling equilibrium	Impulse balance equilibrium	Observed average of 6 observations
Game 7	U	0.091	0.070	0.057	0.056	0.104	0.141
	L	0.909	0.882	0.664	0.691	0.634	0.564
Game 8	U	0.182	0.172	0.185	0.222	0.258	0.250
	L	0.727	0.711	0.619	0.601	0.561	0.586
Game 9	U	0.273	0.250	0.137	0.154	0.188	0.254
	L	0.909	0.898	0.753	0.767	0.764	0.827
Game 10	U	0.364	0.348	0.286	0.308	0.304	0.366
	L	0.818	0.812	0.679	0.731	0.724	0.699
Game 11	U	0.364	0.354	0.286	0.339	0.354	0.331
	L	0.727	0.721	0.679	0.651	0.646	0.652
Game 12	U	0.455	0.449	0.448	0.405	0.466	0.439
	L	0.636	0.634	0.613	0.600	0.604	0.604

comparisons among these three concepts difficult. The cutouts for games 7 to 12 show a similar picture. However, contrary to what happens in other games, in game 9 Nash equilibrium and quantal response equilibrium are slightly nearer to the observed averages than the other three concepts. As we shall see in Section III F, our data suggest that the results of game 9 are influenced by especially large random fluctuations.

As has been explained in Section IB–E, each of the three concepts, Nash equilibrium, action-sampling equilibrium, and quantal response equilibrium, yields the same prediction for the two games in a pair. This is not the case for payoff-sampling equilibrium or impulse balance equilibrium.

B. The Measure of Predictive Success

We look at the five theories compared in this paper as predictions of the relative frequencies of *U* and *L* in an independent subject group playing one of the games 1 to 12. We do not want to assert that a player uses the same mixed strategy in all 200 periods of a session, and we also do not assume that all players in the same role always behave in the same way. Presumably the players are engaged in complex learning processes which differ from person to person. Nevertheless, such behavior may result in frequencies of *U* and *L* that can be predicted reasonably well by stationary concepts. It is important to know how well-observed relative frequencies can be explained without going into the details of stochastic learning models.

For a theory predicting a point in an Euclidian space, the squared distance of theoretical and observed values is a reasonable measure of predictive success, in the sense that the predicted

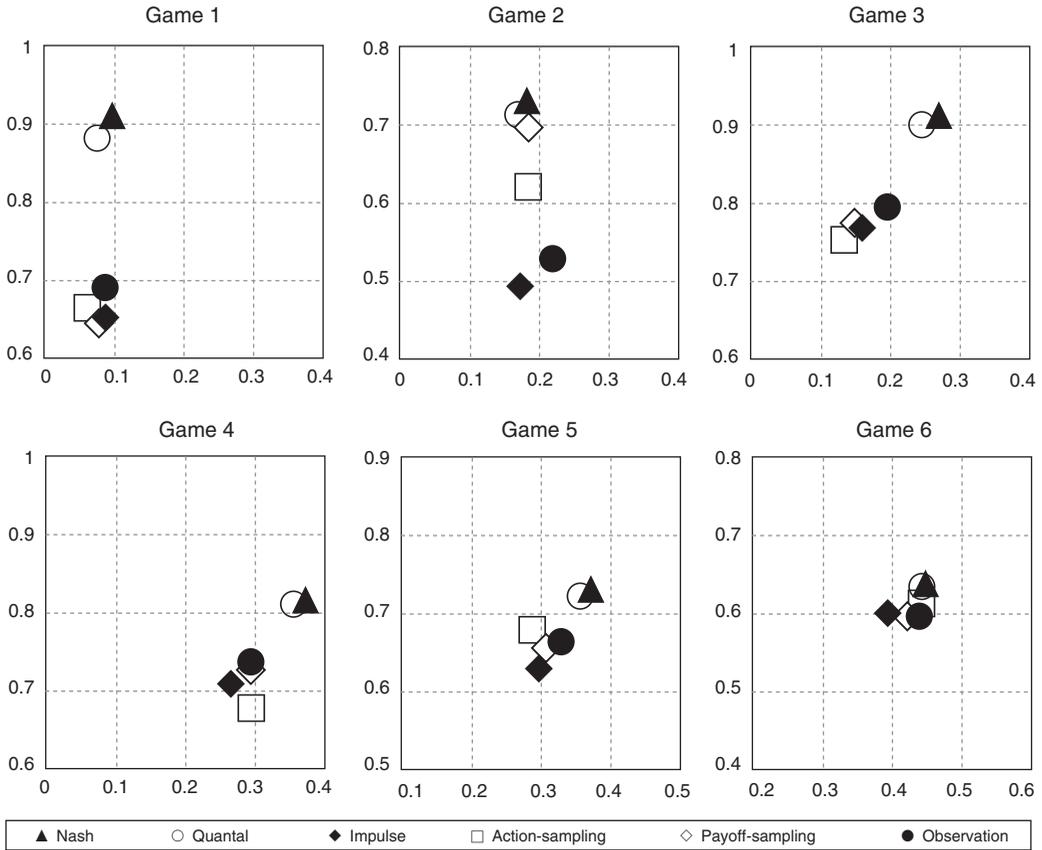


FIGURE 8. VISUALIZATION OF THE THEORETICAL EQUILIBRIA AND THE OBSERVED AVERAGE IN THE CONSTANT SUM GAMES

success is the greater the smaller this distance is. In the following, we explain how this measure is applied to our data. Each game i with $i = 1, \dots, 12$ has been played by s_i independent subject groups with $s_i = 12$ for $i = 1, \dots, 6$ and $s_i = 6$ for $i = 7, \dots, 12$.

We use the index j with $j = 1, \dots, s_i$ for the subject groups. Let f_{iUj} and f_{iLj} be the relative frequencies of U and L in the j -th independent subject group playing game i . Consider a prediction p_U and q_L for these relative frequencies; then,

$$(16) \quad Q_{ij} = (f_{iUj} - p_U)^2 + (f_{iLj} - q_L)^2$$

is the squared distance of the j -th observation for game i from the prediction for game i . The mean squared distance for the data of this game i from (p_U, p_L) is as follows:

$$(17) \quad Q_i = \frac{1}{s_i} \sum_{j=1}^{s_i} Q_{ij}.$$

We shall look at the overall predicted success, but also at the predicted success of the constant sum games 1 to 6 and the nonconstant sum games 7 to 12, separately. Define

$$(18) \quad Q_C = \frac{1}{6} \sum_{i=1}^6 Q_i, \quad Q_N = \frac{1}{6} \sum_{i=7}^{12} Q_i, \quad Q = \frac{1}{12} \sum_{i=1}^{12} Q_i.$$

The indices C and N stand for constant sum and nonconstant sum games. The *mean squared distances* Q_C , Q_N , and Q will be the basis of our comparison of the five theories.

For every game i , let f_{iU} and f_{iL} be the mean values of f_{iUj} and f_{iLj} with $j = 1, \dots, s$:

$$(19) \quad f_{iU} = \frac{1}{s_i} \sum_{j=1}^s f_{iUj} \quad \text{for } i = 1, \dots, 12, \quad f_{iL} = \frac{1}{s_i} \sum_{j=1}^s f_{iLj} \quad \text{for } i = 1, \dots, 12.$$

The expression

$$(20) \quad S_i = \frac{1}{s_i} \sum_{j=1}^s (f_{iUj} - f_{iU})^2 + (f_{iLj} - f_{iL})^2 \quad \text{for } i = 1, \dots, 12$$

is the *sampling variance* of game i and

$$(21) \quad T_i = (f_{iU} - p_U)^2 + (f_{iL} - q_L)^2 \quad \text{for } i = 1, \dots, 12$$

is the *theory specific component* of the mean squared distance. The mean squared distance for a game can be split into these two components:

$$(22) \quad Q_i = S_i + T_i \quad \text{for } i = 1, \dots, 12.$$

Define

$$(23) \quad S_C = \frac{1}{6} \sum_{i=1}^6 S_i, \quad S_N = \frac{1}{6} \sum_{i=7}^{12} S_i, \quad S = \frac{1}{12} \sum_{i=1}^{12} S_i.$$

$$(24) \quad T_C = \frac{1}{6} \sum_{i=1}^6 T_i, \quad T_N = \frac{1}{6} \sum_{i=7}^{12} T_i, \quad T = \frac{1}{12} \sum_{i=1}^{12} T_i.$$

The mean squared distances Q_C , Q_N , and Q can also be split into two components:

$$(25) \quad Q_C = S_C + T_C, \quad Q_N = S_N + T_N, \quad Q = S + T.$$

Note that each game receives equal weight in Q , S , and T , in spite of the fact that there are twice as many observations for each constant sum game than for each nonconstant sum game. This conforms to the goal of obtaining an adequate judgment of the overall goodness of fit for the 12 games.

Since the mean sampling variances S_C , S_N , and S do not depend on the theory under consideration, it does not really matter whether the comparison of theories is based on Q_C , Q_N , and Q or alternatively on T_C , T_N , and T . However, the mean squared distances Q_C , Q_N , and Q are more natural measures of predictive success. A high sampling variance limits the accuracy of prediction, even if the theory-specific component is very small. Therefore, the mean squared distance of the individual observations from the theory is more adequate as a measure of predictive success.

For none of the five theories considered here, the mean squared distance Q can be smaller than S . The sampling variance S is an unavoidable part of Q .

C. Comparison of Sample Sizes for Action-Sampling Equilibrium

Originally, action-sampling equilibrium with the sample size 7 had been considered as a theory to be compared with the data, since this sample size finds some admittedly weak support

in the psychological literature (Miller 1956). The sample size 7 seems to be connected to the average capacity of short-term memory. It is not really clear, however, whether this is relevant for the behavior in our experiments. Therefore, another sample size could have yielded a better fit for our data.

In order to verify this, we compared the predictive success for action-sampling equilibria with different sample sizes.

Figure 9 shows the overall mean squared distances Q for the action-sampling equilibria with the sample sizes $n = 2, \dots, 10$. It can be seen immediately that the average squared distance is smallest for $n = 7$. This means that the best fit to the data is obtained with sample size 7. In our comparison of the five concepts, we therefore do not have to consider other sample sizes for action-sampling equilibrium.

The figure also shows the mean sampling variance in grey. It can be seen that for the sample size 7 the mean squared distance Q is much nearer to its unavoidable part S than for all other sample sizes.

D. Comparison of Sample Sizes for Payoff-Sampling Equilibrium

Figure 10 shows the overall mean squared distances Q for the payoff-sampling equilibria with the sample sizes $n = 1, \dots, 10$. It can be seen that the sample size 6 yields the best fit to the data. Therefore our comparison of the five theories is based on the sample size 6 for payoff-sampling equilibrium.

E. Original versus Transformed Games

The basic idea of impulse balance is applied to the transformed game rather than the original one. This idea could also be applied directly to the original game. As we shall see later, the application to the transformed game yields a better fit to the data. This was already true for the pilot study on games A and B. We therefore decided to test impulse balance theory in the form described in Section IE. It is of interest, however, to examine the question how the direct application compares to the concept of impulse balance equilibrium proposed here.

It could be the case that the predictive power not only of impulse balance equilibrium, but also of other concepts, is increased by applying it to the transformed game rather than to the original one.

We shall examine this question for Nash equilibrium, action-sampling equilibrium, and payoff-sampling equilibrium. Contrary to Nash equilibrium and quantal response equilibrium, action-sampling equilibrium and payoff-sampling equilibrium fit the data quite well. It is therefore of special interest to explore whether a better fit could be obtained by applying these two concepts to the transformed game rather than to the original one. If, in this way, one obtained a better fitting version of one of the two concepts, then this version should be compared with the other three theories.

We did not examine what happens if quantal response equilibrium is applied to the transformed game rather than to the original one. In the cutouts of Figures 8 and 14 in online Appendix A, quantal response equilibrium is always very near to Nash equilibrium and it can be expected that this would not change in an application to the transformed game.

Figure 11 shows the overall mean squared distances for Nash equilibrium, action-sampling equilibrium, payoff-sampling equilibrium, and impulse balance equilibrium applied directly to the original game or to the transformed game. It can be seen that only impulse balance theory profits from being applied to the transformed game, whereas the other three theories do not gain by being modified in this way.

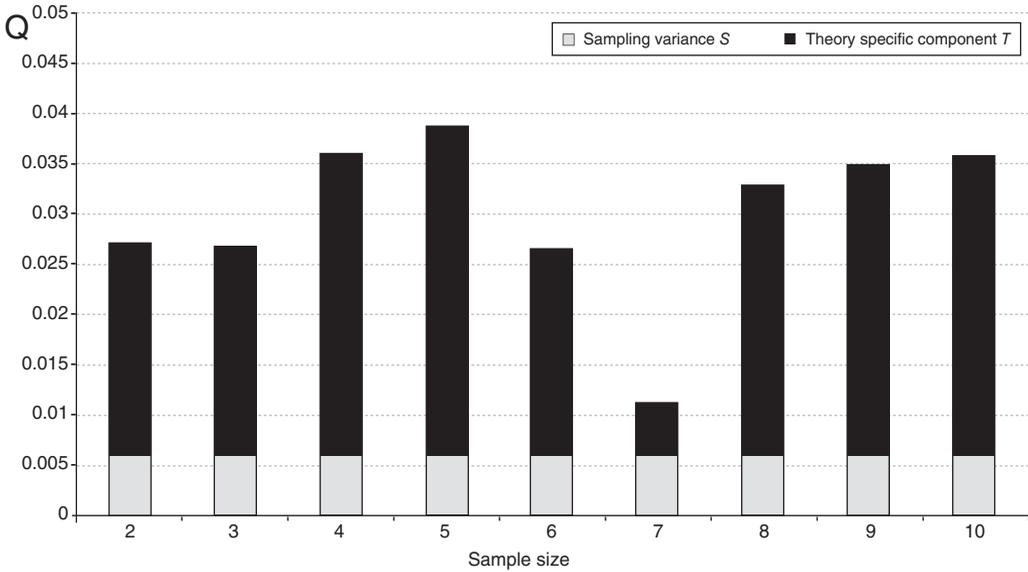


FIGURE 9. OVERALL MEAN SQUARED DISTANCES Q FOR THE ACTION-SAMPLING EQUILIBRIA WITH DIFFERENT SAMPLE SIZES

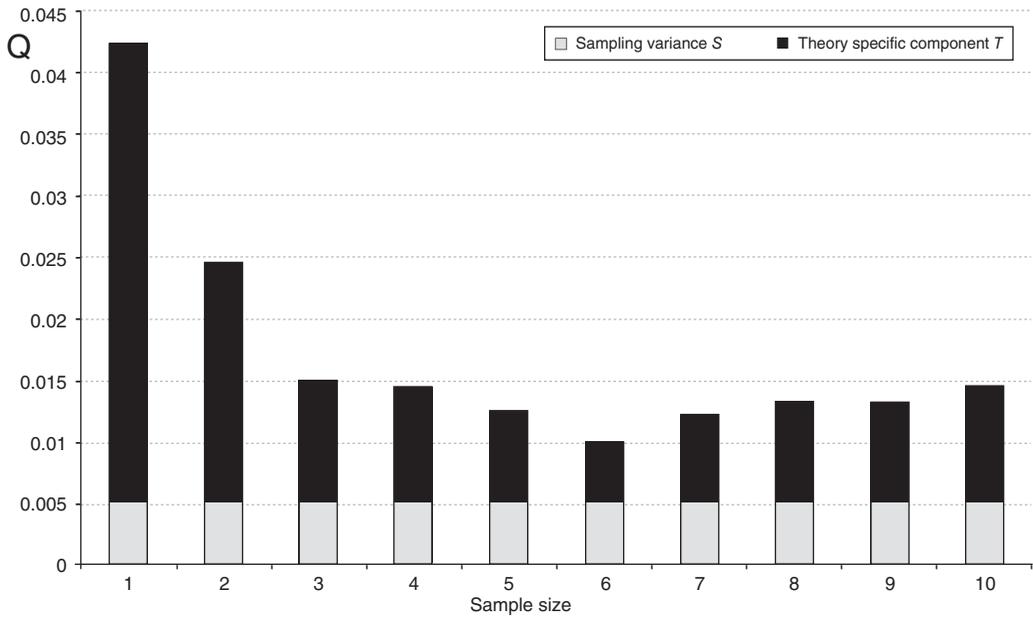


FIGURE 10. OVERALL MEAN SQUARED DISTANCES Q FOR THE PAYOFF-SAMPLING EQUILIBRIA WITH DIFFERENT SAMPLE SIZES

The figure also shows the decomposition of the mean squared distance Q into the sampling variance S (grey) and the theory-specific component T (black and white, respectively). The difference between the applications to the original game and the transformed one are even more dramatic in the case of impulse balance equilibrium if one looks at the theory-specific components instead of the mean squared distance.

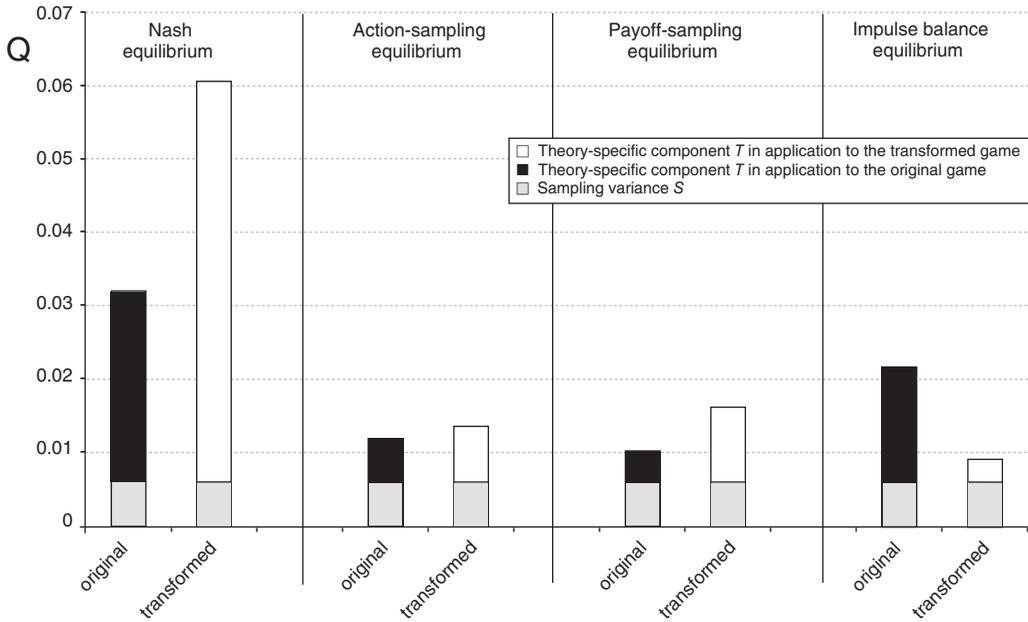


FIGURE 11. ADVANTAGES AND DISADVANTAGES OF APPLYING A CONCEPT TO THE TRANSFORMED GAME RATHER THAN THE ORIGINAL ONE

In view of Figure 7, it seems to be justified not to add the modifications of Nash equilibrium, action-sampling equilibrium, and payoff-sampling equilibrium to the list of the five theories that are the main focus in this paper.

As we shall see in the next section, impulse balance equilibrium fits our data best. Figure 15 (in the online Appendix) shows that this success is not due primarily to the use of the transformed game. Otherwise, the predictive success of other concepts should be improved as well, if they are applied to the transformed game rather than the original one. This is not the case.

F. Comparison of the Five Theories

Table 2 shows the mean squared distances of the 5 theories for the 12 games separately. It also contains the sampling variance for each game.

Figure 12 shows the overall mean squared distances Q for the five theories compared in this paper. It can be seen that there is a clear order of success: impulse balance equilibrium, payoff-sampling equilibrium, action-sampling equilibrium, quantal response equilibrium, and Nash equilibrium. The figure also shows the sampling variance S in grey and the theory-specific components in black.

The sampling variance for game 9 is much greater than for other games. This is probably the reason for the unusual constellation of the cutout for game 9 in Figure 14, online Appendix A.

G. Changes over Time

The question arises whether the order of predictive success of the five theories remains stable over time. Of course we can investigate this question only within the span of the 200 periods played in our experiments. For this purpose, we compared the first hundred periods with the

TABLE 2—SQUARED DISTANCES OF THE FIVE THEORIES

	Nash equilibrium	Quantal response equilibrium	Action- sampling equilibrium	Payoff- sampling equilibrium	Impulse balance equilibrium	Sampling variance
Game 1	0.0572	0.0460	0.0103	0.0112	0.0213	0.00909
Game 2	0.0483	0.0428	0.0164	0.0098	0.0102	0.00693
Game 3	0.0321	0.0250	0.0087	0.0057	0.0073	0.00523
Game 4	0.0169	0.0137	0.0072	0.0041	0.0054	0.00403
Game 5	0.0149	0.0136	0.0115	0.0100	0.0117	0.00953
Game 6	0.0042	0.0039	0.0027	0.0028	0.0045	0.00246
Game 7	0.1237	0.1082	0.0189	0.0253	0.0081	0.00178
Game 8	0.0298	0.0269	0.0106	0.0063	0.0060	0.00531
Game 9	0.0212	0.0192	0.0332	0.0276	0.0224	0.01409
Game 10	0.0208	0.0196	0.0134	0.0109	0.0111	0.00665
Game 11	0.0098	0.0084	0.0059	0.0032	0.0036	0.00307
Game 12	0.0045	0.0042	0.0033	0.0047	0.0039	0.00317

second hundred periods. Figure 13 shows the mean squared distances decomposed into sampling variance (grey) and the theory-specific components (black and white, respectively) for periods 1–100 (left) and 101–200 (right) for the five theories compared in this paper. It can be seen that in the second half of the experiments the predictive success of payoff-sampling equilibrium is slightly greater than that of impulse balance equilibrium. The difference is not significant under the Wilcoxon signed rank test. The predictive success of impulse balance equilibrium is the same one in the first and second half. For each of the other four theories the performance is better in the second than in the first half.

The sampling variance is greater in the second half than in the first half. A two-tailed matched-pairs Wilcoxon signed rank test applied to the sampling variances for the first half and the second half in the 12 games shows no significant difference. Therefore, we interpret the difference between the sampling variances in Figure 13 as due to a random effect.

The improvement of predictive success in the second half of the experiment is connected to a movement of the observed relative frequencies nearer to the convex hull of the theoretical probability vectors. The relative frequencies for the first and the second half of game 4 are both inside the convex hull, but for the other 11 games the relative frequencies for the first half are outside the convex hull. In the second half they are either inside (four games) or still outside but nearer to the convex hull (seven games).

Apart from the reversal between payoff-sampling equilibrium and impulse balance equilibrium, the order of predictive success of the five theories remains unchanged from the first half to the second half of the experiments. The difference between the predictive success of payoff-sampling equilibrium and impulse balance equilibrium in the second half of the experiments is quite small, however, and may be due to random influences. The data do not permit conclusions about convergence to a specific stationary concept over time.

H. Significance of the Comparisons of Predictive Success

In Section IIIA, we pointed out that the discrimination among the five concepts tends to be worse the nearer the games are to the middle of the parameter space. Therefore, we cannot expect significant results for the 12 or 6 observations for each of the games separately. It is more reasonable to apply a test to all constant sum games together, and to do the same for all nonconstant sum games together.

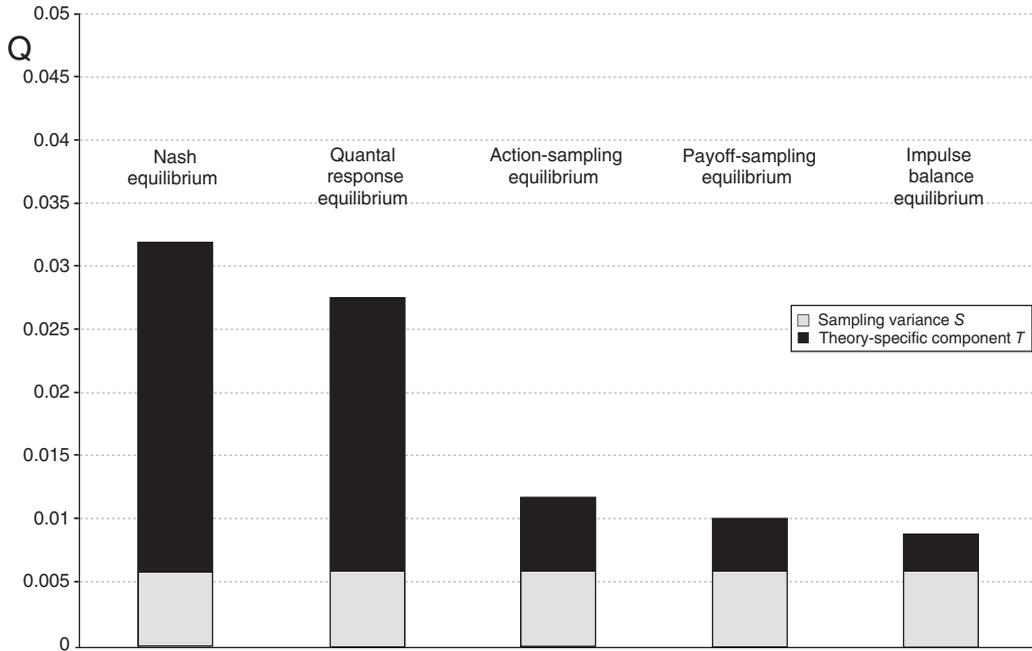


FIGURE 12. OVERALL MEAN SQUARED DISTANCES OF THE FOUR STATIONARY CONCEPTS COMPARED TO THE OBSERVED AVERAGE

In order to compare the performance of two stationary concepts in the 12 games, we apply the Wilcoxon matched-pairs signed rank test to the squared distances of the theoretical values from the observed relative frequencies for the 108 independent subject groups.

In the application of this test, differences of the squared distances are computed for each of the 108 observations and then ranked from 1 to 108 according to their absolute value. Smaller absolute values receive a lower rank. The test statistic is the sum of the ranks in favor of the first theory, in the sense that the squared distance for the first theory is lower than that for the second theory. This means that higher differences count more than lower ones, since they are less likely to be disturbed by random fluctuations. Therefore, the fact that games near the middle of the parameter space discriminate less among the theories is automatically taken into account by the Wilcoxon matched-pairs signed rank test.

The same test has also been applied to the 72 observations on constant sum games and the 36 observations on nonconstant sum games separately.

Table 3 shows the two-tailed significances in favor of the row concept.

In the following, we look at the overall comparisons based on all 108 observations. The comparison between impulse balance equilibrium and payoff-sampling equilibrium is not significant. The same is true for the comparison between payoff-sampling equilibrium and action-sampling equilibrium. This means that the second rank of payoff-sampling equilibrium in the order of predictive success shown in Figure 16 in the online Appendix may be due to random fluctuations. It could be just as well at the first or third place. However, the order among the other four concept is clearly confirmed by the comparisons based on all 108 observations.

We now turn our attention to the comparison between impulse balance equilibrium and payoff-sampling equilibrium for the 72 observations on constant sum games and the 36 observations on nonconstant sum games. Table 3 shows that payoff-sampling equilibrium performs significantly better in constant sum games, whereas impulse balance equilibrium has a significantly greater

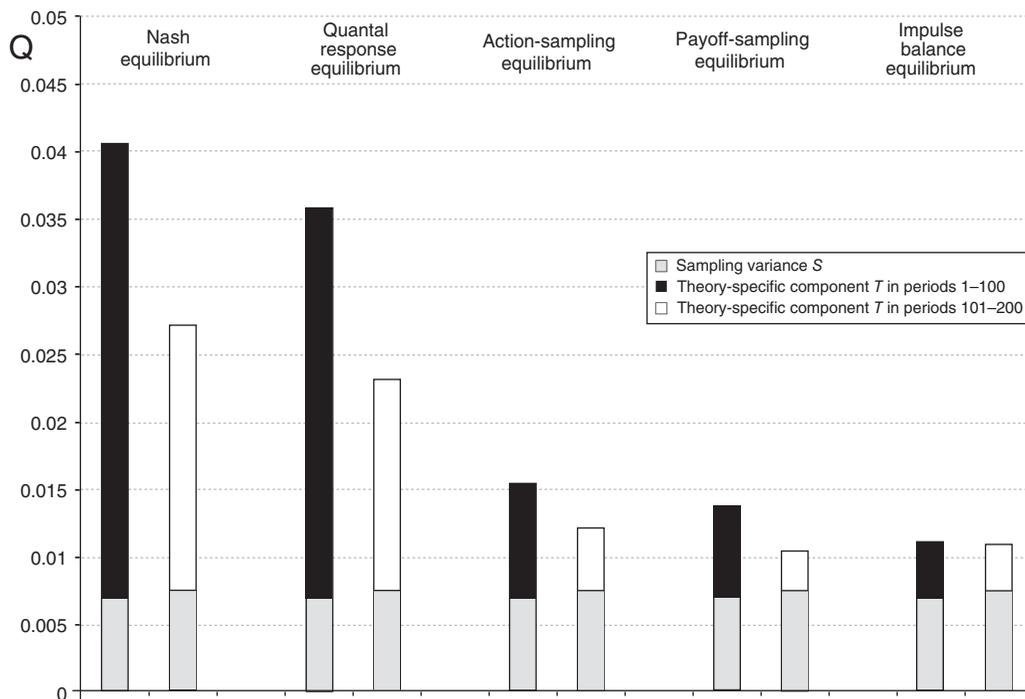


FIGURE 13. COMPARISON OF PREDICTIVE SUCCESS IN THE FIRST HALF AND SECOND HALF OF THE EXPERIMENTS

success in nonconstant sum games. We also see the better fit of payoff-sampling equilibrium for constant sum games and of impulse balance equilibrium for nonconstant sum games in the comparison of both games with action-sampling equilibrium.

It is remarkable that the newer concepts of impulse balance equilibrium, payoff sampling equilibrium, and action-sampling equilibrium clearly outperform the more established concepts of quantal response equilibrium and Nash equilibrium. All the relevant comparisons are highly significant. This is perhaps the most important result of the statistical tests.

According to the Wilcoxon signed rank test, quantal response equilibrium beats Nash equilibrium, in spite of the fact that the predictions of both concepts are very close to each other (see Figures 11 and 12). There is nothing wrong with this. Even if the deviations of predicted and observed values are very similar for both concepts, the relatively small difference between them tends to be in favor of quantal response equilibrium. These differences are important for the Wilcoxon test. Nevertheless, one may say that the relatively small differences are substantially insignificant, even if they are the basis of statistical significance.

IV. Summary and Discussion

Five stationary concepts for completely mixed 2x2-games have been compared in this paper. For this purpose, experiments have been run on 12 games, 6 constant sum games with 12 independent subject groups each, and 6 nonconstant sum games with 6 independent subject groups each.

The games were selected in such a way that the constant sum games were reasonably well distributed over the parameter space. Each nonconstant sum game had the same best reply structure as an associated constant sum game.

TABLE 3—SIGNIFICANCES IN FAVOR OF ROW CONCEPTS, TWO-TAILED MATCHED-PAIRS WILCOXON SIGNED RANK TEST (Rounded to the next higher level among 0.1 percent, 0.2 percent, 0.5 percent, 1 percent, 2 percent, 5 percent, and 10 percent)

	Impulse balance equilibrium	Payoff-sampling equilibrium	Action-sampling equilibrium	Quantal response equilibrium	Nash equilibrium
Impulse balance equilibrium		n.s. — 5 percent	10 percent <i>n.s.</i> 2 percent	1 percent <i>1 percent</i> 1 percent	1 percent <i>1 percent</i> 1 percent
Payoff-sample equilibrium	n.s. <i>10 percent</i> —		n.s. <i>10 percent</i> <i>n.s.</i>	1 percent <i>1 percent</i> 1 percent	1 percent <i>1 percent</i> 1 percent
Action-sample equilibrium				1 percent <i>1 percent</i> 1 percent	1 percent <i>1 percent</i> 0.5 percent
Quantal response equilibrium					1 percent <i>2 percent</i> 5 percent
Nash equilibrium					

Notes: **Above: all 108 Experiments;** *Middle: 72 constant-sum game experiments;* *Below: 36 non-constant sum game experiments.*

Each subject group consisted of eight participants, four playing on one side and four on the other. Each subject group played only one game over 200 periods with random matching.

The literature reports about similar experiments with 2x2-games (McKelvey, Palfrey, and Roberto A. Weber 2000; Goeree, Holt, and Palfrey 2003; Binmore, Swierzbinski, and Proulx 2001; Ochs 1995). Usually, the number of periods played is much lower, and more than one game has been played by the same subjects in one session. Thus, in the experiments by Goeree, Holt, and Palfrey (2003), the number of periods was 40. We wanted a greater number of periods because it is doubtful that a stationary state can be reached within only relatively few periods. Play must be long enough to wash out initial effects.

An exception with respect to the number of plays is the paper by Binmore, Swierzbinski, and Proulx (2001). They report experiments about several games played 150 times. There was only one completely mixed 2x2-game (game 1) among them, however. Each subject played seven games (including two practice games). If several games are played one after the other by the same subjects, transfer of experience may occur from earlier to later games. Moreover, data from different games played by the same subject are not statistically independent from each other. In our experiment each subject participated in only one independent subject group and played only one game. This is necessary for an appropriate application of statistical tests.

In the literature, usually only two of the stationary concepts are confronted with experimental data: Nash equilibrium and quantal response equilibrium. An exception is the paper by Avrami, Güth, and Kareev (2005). They successfully compared impulse balance equilibrium with their data, following the suggestion of one of this paper’s authors (Selten). The new concept of action-sampling equilibrium was never examined before. The same is true for payoff-sampling equilibrium.

Our measure of predictive success forms mean square deviations of observed relative frequencies from predicted probabilities for every game separately, and then takes the average over the 12 games. The comparison of the five theories over the entire time span of 200 periods yields a clear order of predictive success from best to worst:

1. Impulse balance equilibrium
2. Payoff-sampling equilibrium

3. Action-sampling equilibrium
4. Quantal response
5. Nash equilibrium

In order to examine the question to what extent one can exclude the possibility that this ranking of the five concepts is merely the outcome of random fluctuations, we have performed pairwise comparisons with the Wilcoxon matched-pairs signed rank test. The results can be summarized as follows: The rank 2 of payoff-sampling equilibrium is not confirmed. One cannot exclude the possibility that this is due to random fluctuation. The rank of payoff-sampling equilibrium could just as well be 1 or 3. However, the order of the other four concepts is confirmed by significance tests.

A remarkable result can be seen in the fact that the newer concepts of impulse balance equilibrium, payoff-sampling equilibrium, and action-sampling equilibrium clearly outperform the more established concepts of Nash equilibrium and quantal response equilibrium. It can be seen in Figure 12 that impulse balance equilibrium, payoff-sampling equilibrium, and one-sample equilibrium are near to each other with respect to their predictive success. Moreover, the predictive success of these three newer theories is strikingly better than that of the two more established concepts. Moreover, all the statistical comparisons between a newer concept on one side and a more established one on the other side are highly significant in favor of the newer one. This is true not only for comparisons based on all 108 observations, but also for those based on observations for constant sum games or nonconstant sum games only.

The comparisons based on constant sum games and nonconstant sum games alone throw light on the differences between impulse balance equilibrium and payoff-sampling equilibrium. Payoff-sampling equilibrium performs better in constant sum games, whereas impulse balance equilibrium shows the better fit for nonconstant sum games.

The best sample size for payoff-sampling equilibrium was 6. The sample size 7 for action-sampling equilibrium is suggested by the finding that 7 seems to be near to the average number of items that can be kept in short-term memory (Miller 1951). Action-sampling equilibrium with sample size 7 fits our data much better than other sample sizes between 2 and 10.

It is of great importance that even for completely mixed constant sum 2×2 -games, Nash equilibrium and quantal response equilibrium fail in comparison to other concepts.

In this paper, we concentrated on games played repeatedly with random matching by two populations. The literature also reports experiments on 2×2 -games played repeatedly by the same two opponents. Behavior in such games may very well be different from that in games played by populations. If two subjects play the same two-person zero-sum game a hundred times against each other, they will be concerned about not being predictable. This may drive them nearer to maximin strategies. The experimental investigation by B. O'Neill (1987) and an empirical paper by Mark Walker and John Wooders (2001) on "Minimax Play at Wimbledon" suggests that this may be the case.

In our experiments, quantal response equilibrium performs significantly better than Nash equilibrium. Quantal response equilibrium was applied with the same free parameter estimated from the data for all games. This parameter was quite high, probably because it has to accommodate relatively many games with a diverse structure. This is perhaps the reason the predictions of quantal response equilibrium are, on the one hand, very near to those of Nash equilibrium theory and, on the other hand, nevertheless significantly better. The difference between the two theories is statistically significant, even if it is substantially insignificant.

In the literature, much better fits of quantal response equilibrium to observed data are reported (e.g., McKelvey, Palfrey, and Weber 2000). In these studies, however, the noise parameter λ is fitted to each game separately. Since only two probabilities need to be predicted, one for one

strategy of each player, the prediction task is substantially facilitated. Loosely speaking, one may say that estimating a parameter for each game separately does half the job of predicting two numbers.

It would be desirable to complement quantal response equilibrium by a theory that permits the computation of the noise parameter λ as a function of the payoffs of the game. Extended in this way, quantal response equilibrium could become a much more powerful stationary concept. Since at the moment no theory of λ is available, we have applied quantal response theory with λ interpreted as a natural constant, which is the same one for all games.

In the same way as Nash equilibrium and quantal response equilibrium, action-sampling equilibrium is still a concept based on best replies, even if these are not best replies to the equilibrium strategies of the others, but to a random sample of strategies on the other side. Payoff-sampling equilibrium is not based on best replies, but rather on the comparison of samples of payoffs obtained for own choices.

Impulse balance equilibrium is very different from the four other concepts, since it is based neither on best responses nor on payoffs obtained for own choices. Unlike the other four concepts, it cannot be considered to be a modification of Nash equilibrium. Impulse balance is different from optimization, even in one-person decision problems (Selten, Abbink, and Cox 2001; Ockenfels and Selten 2005). Moreover, impulse balance equilibrium is applied to a transformed game. The transformation is based on the idea that losses relative to a natural reference point (the pure strategy maximin payoff) count double.

Impulse balance theory could also be applied to the original game, but the application to the transformed game improves its performance. If Nash equilibrium, action-sampling equilibrium, or payoff-sampling equilibrium is applied to the transformed game rather than the original one, the performance of these concepts becomes worse. The transformation is an important part of impulse balance theory but it is not the only reason for its success.

It is not easy to understand why the predictions of the three newer concepts are not very far apart, in spite of the fact that they are based on very different principles. This is perhaps peculiar to our sample. It would be desirable to devise experiments that permit a better discrimination among the three concepts.

In this paper, we look at stationary concepts without any discussion of learning processes. The comparison of our data with learning processes will be the subject of a later paper. As far as movement over time is concerned we looked only at differences between periods 1–100 and 101–200. We have seen that the order of predictive success of impulse balance theory and payoff-sampling theory reverses from the first half to the second half of the experiments. The reversal is not statistically significant. No other changes of the order of predictive success from the first half to the second half are observed. In the second half, the sampling variance is slightly increased. The predictive success of impulse balance equilibrium is the same in the second half as in the first half, but the other four concepts perform much better in the second half than in the first half. The mean frequencies of individual observations seem to move nearer to the convex hull of the theoretical predictions, even if within a game the variance of the relative frequencies in independent subject groups does not change significantly. One cannot know whether the stationary distribution is reached within the 200 periods, but the evidence conveys the impression that one comes near to it.

Stationary concepts are of great importance, especially if they do not depend on parameters, which have to be adjusted to the data. Impulse balance theory does not involve any such parameters and can be used in theoretical investigations just like Nash equilibrium. It is possible to generalize impulse balance theory to general games in normal form (see online Appendix F). It would certainly be desirable to gain experiences with games with more than two strategies or more than two players.

REFERENCES

- Abbink, Klaus, and Abdolkarim Sadrieh.** 1995. "RatImage—Research Assistance Toolbox for Computer-Aided Human Behavior Experiments." University of Bonn SFB Discussion Paper B-325.
- Avrahami, Judith, Güth Werner, and Yaakov Kareev.** 2005. "Games of Competition in a Stochastic Environment." *Theory and Decision*, 59(4): 255–94.
- Binmore, Ken, Joe Swierzbinski, and Chris Proulx.** 2001. "Does Minimax Work? An Experimental Study." *Economic Journal*, 111(473): 445–64.
- Davis, E. A.** 1931. "Knox Cube Test and Digit Span." *Journal of Genetic Psychology*, 32: 235–37.
- Della Sala, S., C. Gray, A. Baddeley, N. Allamano, and L. Wilson.** 1999. "Pattern Span: A Tool for Unwelding Visuo-Spatial Memory." *Neuropsychologia*, 37(10): 1189–99.
- Erev, Ido, and Alvin E. Roth.** 1998. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed Strategy Equilibria." *American Economic Review*, 88(4): 848–81.
- Goeree, Jakob K., Charles A. Holt, and Thomas R. Palfrey.** 2003. "Risk Averse Behavior in Generalized Matching Pennies Games." *Games and Economic Behavior*, 45(1): 97–113.
- Hart, Sergiu, and Andreu Mas-Colell.** 2000. "A Simple Adaptive Procedure Leading to Correlated Equilibrium." *Econometrica*, 68(5): 1127–50.
- Kahneman, Daniel, Amos Tversky.** 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2): 263–91.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 10(1): 6–38.
- McKelvey, Richard D., Thomas R. Palfrey, and Roberto A. Weber.** 2000. "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2 X 2 Games with Unique Mixed Strategy Equilibria." *Journal of Economic Behavior and Organization*, 42(4): 523–48.
- Miller, George A.** 1956. "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information." *Psychological Review*, 63(2): 81–97.
- Ochs, Jack.** 1995. "Games with Unique, Mixed Strategy Equilibria: An Experimental Study." *Games and Economic Behavior*, 10(1): 202–17.
- Ockenfels, Axel, and Reinhard Selten.** 2005. "Impulse Balance Equilibrium and Feedback in First Price Auctions." *Games and Economic Behavior*, 51(1): 155–70.
- O'Neil, B.** 1987. "Nonmetric Test of the Minmax Theory of Two-Person Zerosum Games." *Proceedings of the National Academy of Sciences*, 84(7): 2106–09.
- Osborne, Martin J., and Ariel Rubinstein.** 1998. "Games with Procedurally Rational Players." *American Economic Review*, 88(4): 834–47.
- Osborne, Martin J., and Ariel Rubinstein.** 2003. "Sampling Equilibrium, with an Application to Strategic Voting." *Games and Economic Behavior*, 45(2): 434–41.
- Selten, Reinhard, and Joachim Buchta.** 1999. "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions." In *Games and Human Behavior: Essays in the Honor of Amnon Rapoport*, ed. David Budescu, Ido Erev, and Rami Zwick, 79–104. Mahwah NJ: Lawrence Associates.
- Selten, Reinhard.** 2000. "Lernrichtungstheorie." In *Innovative Kapitalanlagekonzepte*, ed. Elisabeth Hehn, 11–18. Wiesbaden: Gabler Verlag.
- Selten, Reinhard, Klaus Abbink, and Ricarda Cox.** 2005. "Learning Direction Theory and the Winner's Curse." *Experimental Economics*, 8(1): 5–20.
- Walker, Mark, and John Wooders.** 2001. "Minimax Play at Wimbledon." *American Economic Review*, 91(5): 1521–38.

Copyright of *American Economic Review* is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.