Matrix Games, Mixed Strategies, and Statistical Mechanics

J. Berg* and A. Engel†
Institute for Theoretical Physics, Otto-von-Guericke University, Postfach 4120, D-39016 Magdeburg, Germany
(Received 27 February 1998)

Matrix games constitute a fundamental problem of game theory and describe a situation of two players with completely conflicting interests. We show how methods from statistical mechanics can be used to investigate the statistical properties of optimal mixed strategies of large matrix games with random payoff matrices and derive analytical expressions for the value of the game and the distribution of strategy strengths. In particular the fraction of pure strategies not contributing to the optimal mixed strategy of a player is calculated. Both independently distributed as well as correlated elements of the payoff matrix are considered and the results are compared with numerical simulations. [S0031-9007(98)07803-X]

PACS numbers: 02.50.Le, 05.20.–y, 05.40.+j, 64.60.Cn

Game theory models in mathematical term problems of strategic decision making typically arising in economics, sociology, or international relations and owes much of its modern form to von Neumann [1]. The generic situation in game theory consists of a set of players $\{X, Y, \ldots\}$ choosing between different *strategies* $\{X_i\}, \{Y_i\}, \ldots$, the combination of which determines the outcome of a game specified by the payoffs $P_X(X_i, Y_i, \ldots), P_Y(X_i, Y_i, \ldots), \ldots$ each player is going to receive. The payoffs depend on the strategies of *all* players and the problem for every individual player is to choose *his* strategy such as to optimize his payoff without having control over the strategies of all other players. Despite the extreme simplification of the real world situation inherent in this framework, game theory has proven not only to be a viable mathematical discipline but also to be able to characterize important features of economical systems. Many interesting results have been obtained since von Neumann's pioneering work, including the characterization of equilibria [1,2] and the emergence of cooperation [3]. However, detailed investigations have been restricted either to general statements concerning, e.g., the existence of equilibria, or to situations where every player has only a small number of strategies at his disposal and where the payoffs are simple functions of these strategies. As many situations of interest show a large number of possible strategies and rather complicated relationships between strategic choices and the resulting payoffs, it is tempting to model the payoffs by a random function and to apply the methods of statistical mechanics to describe the properties of the game. This will be a sensible approach if there are characteristic "macroscopic" quantities which do not depend on the particular realization of the random parameters, i.e., are *self-averaging* in the sense of the statistical mechanics of disordered systems [4] (for related applications, see Ref. [5]).

In the present Letter we show how methods from statistical mechanics can be applied to characterize the statistical properties of optimal strategies in matrix games with large randomly chosen payoff matrices. Explicitly, we calculate the mean payoff and the fraction of pure strategies which occur in the optimal mixed strategy of a player. For simplicity we restrict ourselves to matrix games, the type of zero-sum games between two players which also forms the basis of von Neumann's treatment [1,6]. Such games are defined by a (not necessarily square) payoff matrix c_{ij}: Player X may choose between N strategies X_i and player Y between M strategies Y_j, where $i = 1, \ldots, N$ and $j = 1, \ldots, M$. At each step of this game they receive the payoffs $P_X(X_i, Y_j) = -P_Y(X_i, Y_j) =: c_{ij}$. Since player X wishes to gain as large a payoff c_{ij} as possible, whereas player Y must attempt to reach as small a value of c_{ij} in order to maximize his payoff $P_Y(X_i, Y_j) = -c_{ij}$, the goals of the players are completely conflicting. Thus it is appropriate for the players to proceed as follows: Player X knows that when playing strategy X_i he will receive at least the payoff $\min_j c_{ij}$. He therefore chooses strategy X_{i^*}, satisfying $\min_j c_{i^*j} = \max_i \min_j c_{ij}$. Equivalently, player Y plays strategy Y_{j^*} determined by $\max_i c_{ij^*} = \min_j \max_i c_{ij}$ since it minimizes his losses for the optimal choices of X. It is easy to show that $\max_i \min_j c_{ij} \leq \min_j \max_i c_{ij}$ always. The situation is simple if the matrix has a so-called *saddle point*, i.e., if there is a pair i^*, j^* satisfying $\max_i \min_j c_{ij} = c_{i^*j^*} = \min_j \max_i c_{ij}$. In this case, it is optimal for both players to stick to their *pure strategies* X_{i^*} and Y_{i^*}, respectively, since deviations from an optimal strategy by one of the players will lead to a lower payoff for this player. For a large random matrix c the probability for the existence of a such a saddle point vanishes exponentially with the size of the matrix, and the choice of an optimal strategy is less obvious. Since, in this case, $\max_i \min_j c_{ij} < \min_j \max_i c_{ij}$, player X will attempt to achieve a greater gain than his guaranteed minimal gain $\max_i \min_j c_{ij}$ and likewise Y will attempt to achieve a smaller loss than $\min_j \max_i c_{ij}$. To this end they have to prevent their opponent from guessing which strategy they are going to play and choose each strategy with a certain probability x_i and y_j, respectively [1]. A vector x_i of probabilities is called a *mixed strategy* and by the normalization condition is constrained to lie on the N-dimensional simplex. The famous minimax theorem by von Neumann states that for any payoff matrix c there exists a *saddle point of mixed strategies*, i.e., there are two vectors x_i^* and y_j^* such that

$$\max_{\{x_i\}} \min_{\{y_j\}} \sum_{ij} x_i c_{ij} y_j = \sum_{ij} x_i^* c_{ij} y_j^* = \min_{\{y_j\}} \max_{\{x_i\}} \sum_{ij} x_i c_{ij} y_j \,.$$

The expected payoff for the optimal mixed strategies $\nu_c := \sum_{ij} x_i^* c_{ij} y_j^*$ is called the *value of the game* and x_i^*, y_i^* denote optimal mixed strategies of players X and Y since again deviations from an optimal strategy by one of the players will lead to a lower payoff for this player.

In the following we show how the statistical properties of such optimal mixed strategies for random payoff matrices may be characterized analytically in the limit $N \to \infty, M \to \infty$ with $M/N = \alpha = O(1)$. As a result of the central limit theorem, only the first two cumulants of the probability distribution $P(\{c_{ij}\})$ are relevant, as is generally the case in fully connected disordered systems described by mean-field theories. Since an average value $\langle\langle c \rangle\rangle$ of the elements of the payoff matrix results only in a modified value of the game $\nu_c + \langle\langle c \rangle\rangle$ without changing the optimal mixed strategies, we may set $\langle\langle c \rangle\rangle = 0$ without loss of generality and take the elements c_{ij} to be independent Gaussian distributed variables with zero mean and variance N^{-1}.

We then note [6] that a necessary and sufficient condition for the mixed strategy $\{x_i\}$ of player X to be optimal is

$$\sum_i x_i c_{ij} \geq \nu_c \;\forall\; j \,. \tag{1}$$

The condition is necessary since if violated for some j player Y playing Y_j will lead to a payoff lower than ν_c. It is also sufficient since combining (1) with the minimax theorem gives $\sum_{ij} x_i c_{ij} y_j^* = \nu_c$. We may thus characterize mixed strategies of player X by introducing the partition function

$$Z(\nu) = \frac{\prod_{i=1}^N (\int_0^\infty dx_i) \delta(\sum_{i=1}^N x_i - N) \prod_{j=1}^{\alpha N} \Theta(\sum_i x_i c_{ij} - \nu)}{\prod_{i=1}^N (\int_0^\infty dx_i) \delta(\sum_{i=1}^N x_i - N)}\,, \tag{2}$$

where $\Theta(x)$ is the Heaviside step function and the probabilities of playing a given strategy and the payoff have been rescaled so that $\sum_{i=1}^N x_i = N$. Thus $Z(\nu)$ equals the fraction of the simplex obeying $\sum_i x_i c_{ij} \geq \nu \;\forall\; j$ and therefore lies on the interval $[0, 1]$. Since $Z(\nu)$ scales exponentially with N, the quantity central to our calculation is the entropy $S(\nu) := 1/N \ln Z(\nu)$, which thus takes on values between $-\infty$ and zero as usual for classical systems with continuous degrees of freedom.

Assuming the entropy $S(\nu)$ to be self-averaging, we use the replica trick $\ln Z = \lim_{n\to 0} \frac{d}{dn} Z^n$ and compute the average over the payoffs of the replicated partition function for integer n $(a, b = 1, \ldots, n)$. The calculation proceeds by using the integral representation of the Heaviside step function and by introducing the symmetric matrix of overlap order parameters $q_{ab} = 1/N \sum_i x_i^a x_i^b$ via integrals over q_{ab} and delta functions represented by integrals over the conjugate order parameters $\hat{q}_{ab}$ [7]. The integrals over E_a arise from the integral representation of the constraint $\sum_i x_i^a = N$ giving

$$\begin{aligned}\langle\langle Z^n(\nu) \rangle\rangle = &\prod_{a \geq b} \int \frac{dq_{ab}\, d\hat{q}_{ab}}{2\pi/N} \prod_a \int \frac{dE_a}{2\pi/N} \exp\left(-iN \sum_{a\geq b} q_{ab}\hat{q}_{ab} - iN \sum_a E_a - N \sum_a 1 \right) \\ &\times \prod_{a,i} \int_0^\infty dx_i^a \exp\left(i \sum_{a \geq b, i} \hat{q}_{ab} x_i^a x_i^b + i \sum_{a,i} E_a x_i^a \right) \\ &\times \prod_{a,j} \int_\nu^\infty d\lambda_j^a \int \frac{dy_j^a}{2\pi} \exp\left(-\frac{1}{2} \sum_{a,b,j} q_{ab} y_j^a y_j^b + i \sum_{a,j} y_j^a \lambda_j^a \right).\end{aligned} \tag{3}$$

In the limit of large payoff matrices $N \to \infty$ the integrals over order parameters are dominated by their saddle point. Throughout this paper we use the replica-symmetric ansatz [8]

$$\begin{aligned} q_{aa} &= q_1 \qquad i\hat{q}_{aa} = -1/2\hat{q}_1 \qquad iE_a = E \;\forall\; a \\ q_{ab} &= q_0 \qquad i\hat{q}_{ab} = \hat{q}_0 \qquad\qquad\quad \forall a > b\,. \end{aligned} \tag{4}$$

The limit $n \to 0$ of (3) may now be taken by analytic continuation giving an entropy

$$\begin{aligned} S(\nu) = \text{extremum}_{q_1,q_0,E,\hat{q}_0,\hat{q}_0} \Bigg[&\frac{1}{2} q_1 \hat{q}_1 + \frac{1}{2} q_0 \hat{q}_0 - E + \frac{1}{2} \ln(2\pi) + \alpha \int \mathrm{D}s \ln H\left(\frac{\sqrt{q_0}\, s + \nu}{\sqrt{q_1 - q_0}} \right) \\ &- \frac{1}{2} \ln(\hat{q}_1 + \hat{q}_0) - 1 + \frac{\hat{q}_0 + E^2}{2(\hat{q}_1 + \hat{q}_0)} + \int \mathrm{D}r \ln H\left(-\frac{\sqrt{\hat{q}_0}\, r + E}{\sqrt{\hat{q}_1 + \hat{q}_0}} \right) \Bigg], \end{aligned} \tag{5}$$

where $\mathrm{D}s = \frac{ds}{\sqrt{2\pi}} \exp(-s^2/2)$, and $H(x) = \int_x^\infty \mathrm{D}s$ (accordingly for Dr). Numerical evaluation of (5) shows that $S(\nu)$ is a continuously decreasing function of ν. At ν_c it tends to $-\infty$, indicating that, for larger values of ν there are no more solutions to (1). Furthermore as $\nu \to \nu_c$ one finds $q_0 \to q_1$ indicating that as the points contributing to (2) crowd into an ever decreasing area of the simplex which shrinks to a point at ν_c, their mutual overlap q_0 approaches the self-overlap q_1.

In this regime the entropy may be conveniently written in terms of the order parameters $q_0, \hat{q}_0, E$, $\hat{w} = \hat{q}_1 + \hat{q}_0$ and $v = q_1 - q_0$. For $\nu < \nu_c$, $S(\nu)$ describes suboptimal strategies. As $v \to 0$ we find $\hat{q}_0 \sim v^{-2}, \hat{w} \sim v^{-1}$. Rescaling the conjugate order parameters accordingly and expanding the saddle-point equations to leading order in v as $v \to 0$, we find

$$\begin{aligned} \hat{w} - \alpha H(-\nu_c/\sqrt{q_0}) &= 0\,, \\ \hat{w} - H(-E/\sqrt{\hat{q}_0}) &= 0\,, \\ \hat{q}_0 - (\nu_c^2 + q_0)\hat{w} - \alpha\sqrt{q_0}\,\nu_c G(-\nu_c/\sqrt{q_0}) &= 0\,, \\ q_0 - (E^2 + \hat{q}_0)/\hat{w} - \sqrt{\hat{q}_0}\,E/\hat{w}^2 G(-E/\sqrt{\hat{q}_0}) &= 0\,, \end{aligned} \tag{6}$$

with $E = q_0\hat{w} - \hat{q}_0$.

The statistical properties of optimal strategies $\{x_i^*\}$ may be deduced from the proportion of strategies X_i with $x_i^* > a$,

$$\theta(a) := \left\langle\!\!\left\langle (1/N)\sum_i \Theta(x_i^* - a) \right\rangle\!\!\right\rangle = H\!\left(\frac{\hat{w}a - E}{\sqrt{\hat{q}_0}}\right). \tag{7}$$

Thus only a fraction $\theta(0) = \hat{w}$ of the pure strategies X_i have $x_i > 0$ and are played with nonzero probability. This striking effect may be explained by considering the behavior of player Y, whose optimal mixed strategy y_j^* obeys $\lambda_i^* = \sum_j c_{ij} y_j^* \le \nu_c\ \forall\ i$. Since $\nu_c = (1/N)\sum_i x_i^* \lambda_i^*$, x_i^* must be zero if $\lambda_i^* < \nu_c$. This mechanism thus ensures an expected payoff ν_c to X, even if Y chooses an optimal strategy. However it is not to be confused with the concept of domination, widely discussed in the game theory literature [1,6,9], where a strategy X_i has $x_i = 0$ because whatever the response of the opponent some other pure or mixed strategy will lead to a higher expected payoff. In fact, in the thermodynamic limit, domination of a pure strategy occurs with probability zero since for a pure strategy X_k to be dominated by a mixed strategy x_i^D requires $(1/N)\sum_i x_i^D c_{ij} \ge c_{kj}\ \forall\ j$ but the left-hand side is $O(N^{-1})$ whereas the right-hand side is $O(N^{-1/2})$.

Figure 1 shows the value of the game and (inset) the fraction of strategies played with nonzero probability as a function of the aspect ratio α of the payoff matrix. At $\alpha = 1$, $\nu_c = 0$ and $\theta(0) = 1/2$. The result $\nu_c = 0$ at $\alpha = 1$ is a consequence of the symmetry of the distribution of payoffs under $c_{ij} \to -c_{ji}$, i.e., under the interchange of player X and player Y [10]. For $\alpha > 1$, player Y has a greater choice of strategies than player X and vice versa. As expected, the payoff to player X decreases as the range of strategy choices of player Y increases. The fraction of strategies played with nonzero probability increases with α, which reflects the decrease of ν_c with α: At lower ν_c there are fewer i with $\lambda_i^* = \sum_j c_{ij} y_j^* < \nu_c$, so as argued above the number of strategies X_i played with nonzero probability increases as a result.

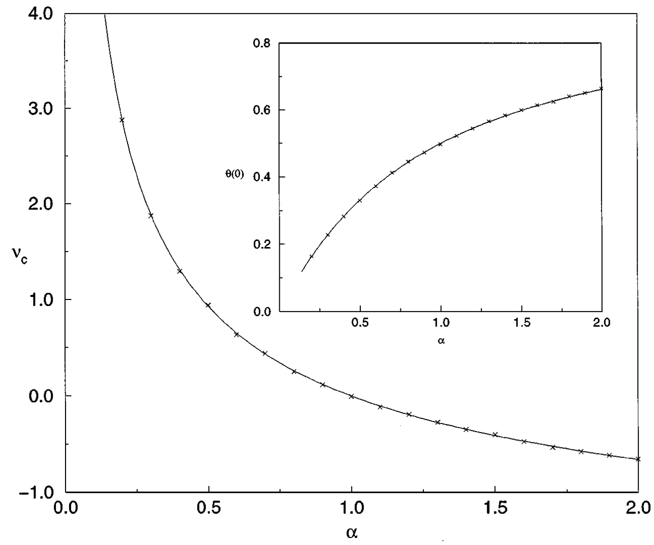

FIG. 1. The value of the game ν_c and (inset) the fraction of strategies played with nonzero probability $\theta(0)$ as a function of α. The analytical results (full line) are compared to numerical simulations with $N = 200$ averaged over 200 samples. The symbol size corresponds to the statistical error.

We next abandon the initial assumption that the individual entries c_{ij} in the payoff matrix are independently distributed and consider the case where the outcomes of the game for different strategy choices of the players are correlated with each other. Such correlations may arise quite naturally in real applications since we expect some strategies to have broadly similar properties and hence yield similar results for a given response of the respective opponent. For simplicity we restrict the discussion to the case $\alpha = 1$. The most general tractable case appears to be

$$\langle\langle c_{ij} c_{kl} \rangle\rangle / (\langle\langle c_{ij} \rangle\rangle \langle\langle c_{kl} \rangle\rangle) =: C_{(ij)(kl)} = C_{ik}^c C_{jl}^r\,, \tag{8}$$

where C_{ik}^c and C_{jl}^r refer to columnlike and rowlike correlations. Of course $P(\{c_{ij}\})$ is not uniquely determined by its second moments, but as argued above it suffices to consider Gaussian distributed payoff matrices. The specific form of C_{ik}^c and C_{jl}^r that we will consider in the following is

$$C_{ik}^{c,r} = \begin{cases} 1 & i = k \\ c_{c,r}/N & i \neq k \end{cases}, \tag{9}$$

and the resulting replica symmetric entropy may be calculated as outlined for the case of uncorrelated payoffs above. Again in the limit $v \to 0$, the corresponding saddle-point equations describe optimal strategies. For $c_c = c_r = c$ the optimal payoff is zero, as a result of the symmetry of $P(\{c_{ij}\})$ under the exchange of players.

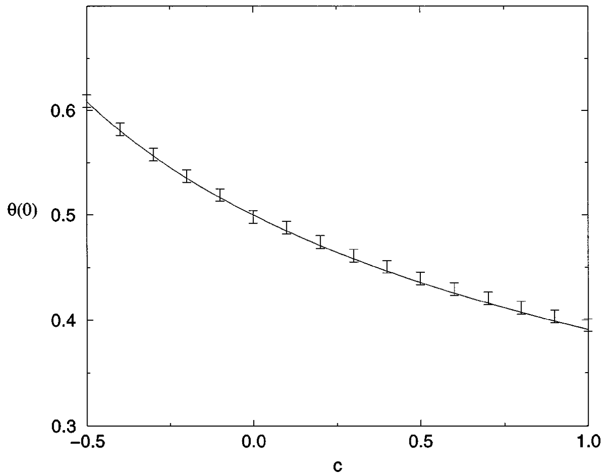

FIG. 2. The fraction $\theta(0)$ strategies played with nonzero probability as a function of $c = c_r = c_c$.

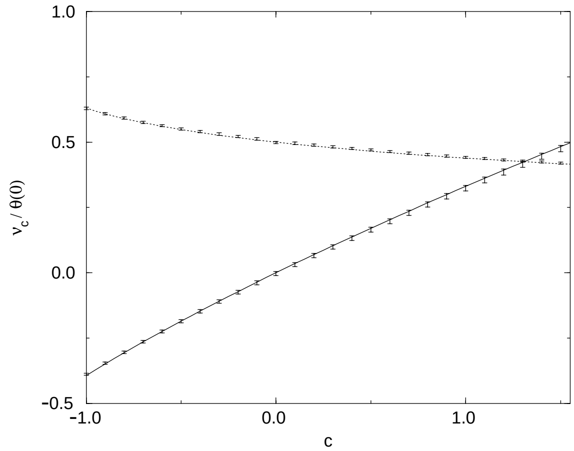

FIG. 3. The optimal payoff ν_c (full) and the fraction $\theta(0)$ of strategies played with nonzero probability (dotted) against $c = c_r$ at $c_c = 0$.

Figure 2 shows the fraction $\theta(0)$ of strategies played with nonzero probability in optimal strategies as a function of c. $\theta(0)$ decreases with increasing c: At positive c there are strategies which tend to be beneficial for player X whatever the response of the opponent. As a result, X concentrates on a smaller fraction of his strategies and vice versa for negative c.

In the asymmetric case $c_r = c, c_c = 0$, however, a nonzero value of the game is possible. The resulting value for ν_c and the fraction of strategies played with nonzero probability are shown in Fig. 3. Again, positive correlations between payoffs in the same row of the payoff matrix lead to strategies which tend to be either beneficial or detrimental to player X. By admitting only the beneficial ones into his mixed strategies, X may achieve a positive payoff. The fraction of strategies played with nonzero probability decreases accordingly.

The simulation results shown in Figs. 1–3 were obtained using the simplex algorithm to solve the linear programming problem [6] defined by (1) for a system of size $N = 200$ averaged over 200 payoff matrices with Gaussian distributed inputs [11]. The numerical results show very good agreement with the analytical expressions.

In conclusion, we have shown that techniques from the statistical mechanics of disordered systems may be used to analyze the statistical properties of optimal solutions of matrix games with random payoffs. Self-averaging macroscopic quantities such as the value of the game were identified and calculated for various probability distributions. These quantities include the fraction of strategies played with nonzero probability. Further problems in matrix games which may be treated using these methods include the effects of deviating from the optimal strategy and the influence of perturbations of the payoff matrix on the optimal strategy, which form the basis of the justification for the full stability of mixed equilibria [12]. Furthermore, work is in progress on the statistical description of Nash equilibria in bimatrix games.

We thank R. Monasson, L. Schneider, and M. Weigt for stimulating discussions. J. B. gratefully acknowledges financial support from the *Studienstiftung des Deutschen Volkes*.

*Email address: johannes.berg@physik.uni-magdeburg.de
†Email address: andreas.engel@physik.uni-magdeburg.de

[1] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton Press, Princeton, NJ, 1953), 3rd ed.

[2] J. F. Nash, Ann. Math. **54(2)**, 286 (1951).

[3] R. Axelrod, *The Evolution of Cooperation* (Basic Books, New York, 1984).

[4] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).

[5] Y. Fu and P. W. Anderson, J. Phys. A **19**, 1605 (1986); S. Diederich and M. Opper, Phys. Rev. A **39**, 4333 (1989); R. Monasson and R. Zecchina, Phys. Rev. Lett. **76**, 3881 (1996); S. Galluccio, J.-P. Bouchaud, and M. Potters, cond-mat/9801209, A. Crisanti *et. al., Products of Random Matrices in Statistical Physics* (Springer, Berlin, 1993).

[6] W. Jianhua, *The Theory of Games* (Oxford University Press, New York, 1988).

[7] E. Gardner, J. Phys. A **21**, 257 (1988).

[8] Replica symmetry may be justified from the convexity of the solution space of Eq. (1); furthermore, it has been found to be locally stable for $\nu < \nu_c$.

[9] A. Brandenburger, J. of Econ. Perspect. **6**, 83 (1992).

[10] Under $\alpha \to 1/\alpha$, $\nu_c \to -\sqrt{\alpha}\,\nu_c$ since the variance of the payoffs is N^{-1}.

[11] Virtually identical results have been obtained for different probability ensembles such as $c_{ij} = \pm N^{-1/2}$.

[12] J. C. Harsanyi, Int. J. Game Theory **2**, 1 (1973).