Noname manuscript No.
(will be inserted by the editor)

Nash equilibria of Cauchy-random zero-sum and coordination matrix games

David P. Roberts

Division of Science and Mathematics, University of Minnesota Morris, Morris, MN 56267 e-mail: `roberts@morris.umn.edu`

The date of receipt and acceptance will be inserted by the editor

Abstract We consider zero-sum games $(A, -A)$ and coordination games (A, A), where A is an m-by-n matrix with entries chosen independently with respect to the Cauchy distribution. In each case, we give an exact formula for the expected number of Nash equilibria with a given support size and payoffs in a given range, and also asymptotic simplications for matrices of a fixed shape and increasing size. We carefully compare our results with recent results of McLennan and Berg on Gaussian random bimatrix games (A, B), and describe how the three situations together shed light on random bimatrix games in general.

Key words Nash equilibrium; support size; Cauchy distribution; zero-sum game; coordination game

1 Introduction

We work throughout in the classical setting of m-by-n bimatrix games (A, B). It is natural to ask about the number of Nash equilibria and the distribution of their support size k and payoffs (λ_1, λ_2) for "random" games. Thrall and Falk (1965) and Faris and Maier (1987) pursued such questions in the zero-sum setting $(A, -A)$. They encountered intractable multivariate integrals and focused instead on numerical results. Various asymptotic statements were heuristically derived in Berg and Engel (1998), Berg and Weigt (1999), and Berg (2000) via remarkable statistical mechanics techniques. Finally, a definitive result when A and B are independent was recently obtained by McLennan and Berg (2005).

Our Theorems 1 and 2 treat two extreme cases of bimatrix games, obtaining a simple formula in each case. The key technical idea is to make randomness rigorous by using the Cauchy density

$$f_c(u) = \frac{1}{\pi}\frac{1}{u^2+1}, \tag{1}$$

instead of the uniform or Gaussian density used by the above authors. Theorem 1 then gives the exact joint distribution of $(k, \lambda_1, \lambda_2) = (k, \lambda, -\lambda)$ in the case of zero-sum games $(A, -A)$. Similarly, Theorem 2 gives the exact joint distribution of $(k, \lambda_1, \lambda_2) = (k, \lambda, \lambda)$ in the case of coordination games (A, A). The zero-sum case is special in that it is known *a priori* that a random game has only one Nash equilibrium. This fact is used crucially in the proof of Theorem 1. In turn, Theorem 1 is used in the proof of Theorem 2. Figure 1 and its caption give a first indication of our results in the special case of square matrices, where $\kappa_1 = k/m = k/n$ is the common support fraction.

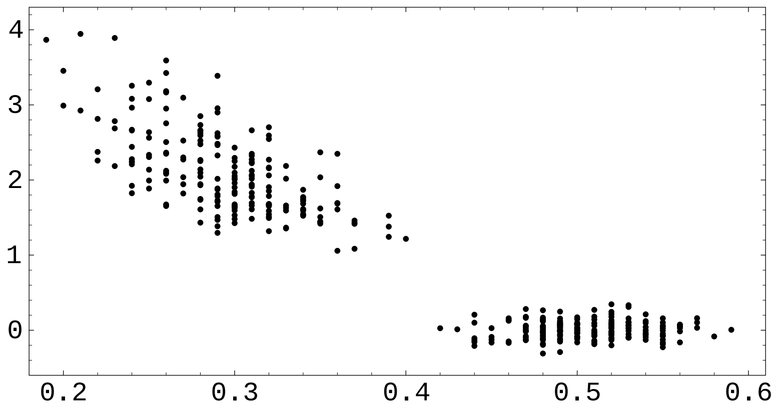

Fig. 1 Pairs (κ_1, λ) corresponding to 100-by-100 Cauchy-random games. The 200 points in the bottom right represent the unique equilibria in 200 zero-sum games. A 100-by-100 Cauchy-random coordination game has on average about 2.66×10^{16} Nash equilibria. The 200 points in the top left represent 200 randomly chosen equilibria from 100-by-100 coordination games. If 100 were replaced by $100u^2$, then the corresponding clouds would contract by a linear factor of u about their respective asymptotic means, $(1/2, 0)$ and $(1 - 2^{-1/2}, \tan(2^{-3/2}\pi)) \approx (0.29, 2.02)$.

Section 2 states our theorems precisely and places them in the context of previous work on Nash equilibria of m-by-n games, especially the McLennan-Berg theorem which is in a strong sense intermediate between our two theorems. Section 3 presents a lemma on Cauchy-random bimatrix games, valid for general correlations, not just our extremes $B = \pm A$. Section 4 proves Theorems 1 and 2. Section 5 centers on Corollaries 1 and 2, each of which describes asymptotic behavior associated to the corresponding theorem. Sections 3-5 all close by elaborating on the relation of our work with the McLennan-Berg theorem, with (42)-(46) at the end of Section 5 refining the McLennan-Berg asymptotics.

Section 6 concludes by graphing and discussing the main asymptotic quantities. We regard our Theorem 1, the McLennan-Berg theorem, and our Theorem 2 as three quantitative anchor points. We describe how these three points give one confidence in certain qualitative expectations about Nash equilibria of random bimatrix games for general notions of randomness, or even in situations where "random" is not mathematically well-defined.

2 Statement of theorems

First we review the setting of bimatrix games and set up our notations. There are two players and a game between them is given by a pair of m-by-n matrices (A, B). Players 1 and 2 move simultaneously, choosing probability vectors $x = (x_1, \dots, x_m)$ and $y = (y_1, \dots, y_n)^t$ respectively. The game ends by Player 1 and 2 receiving payments $\lambda_1 = xAy$ and $\lambda_2 = xBy$ respectively.

A pair of strategies (x, y) is called a *Nash equilibrium* if neither player can improve his return by making a unilateral change in strategy. In general, there may be infinitely many Nash equilibria. However outside a subset of codimension one in matrix space, games (A, B) have only finitely many Nash equilibria. In this case, the strategies x and y in a Nash equilibrium (x, y) have the same number $k \in \{1, \dots, \min(m, n)\}$ of non-zero components. We call k the *support size* of (x, y).

Fix a probability measure measure μ with a continuous density f_μ. In considering random bimatrix games (A, B), we always require that the entries a_{ij} and b_{ij} be all distributed with respect to μ. We allow correlation between a_{ij} and b_{ij}, but otherwise require independence among the entries. We mostly restrict attention to the three simplest types of correlation, which we index by $t = -$, 0, and $+$. In order, these are the zero-sum case $a_{ij} = -b_{ij}$, the neutral case where a_{ij} and b_{ij} are chosen independently, and the coordination case $a_{ij} = b_{ij}$. The main problem is then to compute $E^{\mu,t}_{m,n,k}$, the average number of Nash equilibria of support size k for μ-random m-by-n t-correlated bimatrix games.

For $k = 1$ and all μ, the answer is given by three easily derived formulas,

$$E^{\mu,-}_{m,n,1} = \frac{m!n!}{(m+n-1)!}, \tag{2}$$

$$E^{\mu,0}_{m,n,1} = 1, \tag{3}$$

$$E^{\mu,+}_{m,n,1} = \frac{mn}{m+n-1}. \tag{4}$$

Formula (2) first appeared in Goldman (1957). Formulas (3) and (4) have been the starting point for recent research, see e.g. Powers (1990), Stanford (1995) and Stanford (1999), Roberts (2005) respectively. However the case $k = 1$ is not representative of the general case because it is only in this case that $E^{\mu,t}_{m,n,k}$ is independent of μ.

Thrall and Falk (1965) and Faris and Maier (1987) sought formulas for $E^{\mu,-}_{m,n,k}$ for μ a uniform or Gaussian measure and $k > 1$ and obtained numerical results. Their work provided the starting point for our work here.

We work instead with the Cauchy measure c with density function (1) and cumulative distribution function $F_c(u) = 1/2 + \arctan(u)/\pi$. Then the integrals involved become feasible, and give information about not just support sizes but also payoffs.

Theorem 1 *For Cauchy-random m-by-n zero-sum games, the probability that the unique Nash equilibrium has support size k and payoff* $(\lambda, -\lambda)$ *with* $a \leq \lambda \leq b$ *is*

$$\binom{m}{k}\binom{n}{k} k \int_a^b F_c(\lambda)^{m-1} F_c(-\lambda)^{n-1} f_c(\lambda) d\lambda. \tag{5}$$

For $(a, b) = (-\infty, \infty)$, the integral evaluates to the Beta-value $B(m, n) = (m-1)!(n-1)!/(m+n-1)!$. Thus

$$E^{c,-}_{m,n,k} = \binom{m}{k}\binom{n}{k} kB(m, n), \tag{6}$$

a generalization of (2) in the setting $\mu = c$. Theorem 1 says in particular that k and λ are statistically independent, as only k appears outside the integral in (5) while only λ appears inside the integral. Figure 1 illustrates this independence, and contrasts it with the negative correlation one has in the setting of Theorem 2.

McLennan (2005) worked with the Gaussian measure g with density $f_g(u) = \exp(-u^2/2)/\sqrt{2\pi}$ and cumulative distribution function $F_g(u) = 1/2 + \mathrm{erf}(u/\sqrt{2})/2$. He obtained a result in the neutral case $t = 0$ even for games with arbitrarily many players. McLennan and Berg (2005) made this result more explicit in the case of two players and derived asymptotic consequences. The main two-player result gives simultaneous information about support sizes and renormalized payoffs $\Lambda_1 = \lambda_1/||y||$ and $\Lambda_2 = \lambda_2/||x||$, analogously to (5). Here $|| \cdot ||$ indicates Euclidean norm, as in $||(x_1, x_2)|| = \sqrt{x_1^2 + x_2^2}$. Integrating over all Λ_1 and Λ_2, one gets

$$\begin{aligned} E^{g,0}_{m,n,k} =& \binom{m}{k}\binom{n}{k} \frac{k}{2^{2k-2}\pi} \cdot \\ & \int_{-\infty}^{\infty} F_g(\Lambda_1)^{m-k} e^{-k\Lambda_1^2/2} d\Lambda_1 \int_{-\infty}^{\infty} F_g(\Lambda_2)^{n-k} e^{-k\Lambda_2^2/2} d\Lambda_2, \end{aligned} \tag{7}$$

a generalization of (3) in the setting $\mu = g$.

The referee of the first version of this paper pointed us to the more recent work on this subject. Inspired by the McLennan-Berg theorem and its formal similarity to our Theorem 1, we then found the following theorem for $t = +$. Again for us, the good measure to use is the Cauchy measure c.

Theorem 2 *For Cauchy-random m-by-n coordination games, the expected number of Nash equilibria with support size k and payoff* (λ, λ) *satisfying* $a \leq \lambda \leq b$ *is*

$$\binom{m}{k}\binom{n}{k} k \int_a^b F_c(\lambda)^{m+n-k-1} F_c(-\lambda)^{k-1} f_c(\lambda) d\lambda. \tag{8}$$

For $(a,b) = (-\infty, \infty)$, this becomes

$$E^{c,+}_{m,n,k} = \binom{m}{k}\binom{n}{k} kB(m+n-k,k), \tag{9}$$

a generalization of (4) in the case $\mu = c$.

3 A lemma on the persistence of equilibria in the presence of additional pure strategies

From the definition of Nash equilibrium, it is clear that (x,y) is an equilibrium for an m-by-n bimatrix game (A,B) if and only if

$$x'Ay \leq xAy \text{ for all } x' \in \Delta^r_m, \tag{10}$$

$$xBy' \leq xBy \text{ for all } y' \in \Delta^c_n. \tag{11}$$

Here Δ^r_m is the simplex of probability row m-vectors and Δ^c_n is the simplex of probability column n-vectors.

In the proof of Lemma 1, we use that convolution is additive with respect to width for the Cauchy measure. To say this precisely, recall that for probability measures μ_1 and μ_2 with densities $f_1(u)$ and $f_2(u)$ respectively, the convolution $\mu_1 * \mu_2$ is the probability measure with density $\int_{-\infty}^{\infty} f_1(x)f_2(u-x)dx$. For $w > 0$, let c_w be the width w Cauchy measure with density

$$f_{c_w}(u) = \frac{1}{w} f_c(u/w) = \frac{1}{\pi}\frac{w}{w^2+u^2}.$$

Then the precise statement is $c_{w_1} * c_{w_2} = c_{w_1+w_2}$. For detailed elementary proofs, see either Dwass (1985) or Nelson (1985); for the general context of stable distributions, see Rose and Smith (2002).

Lemma 1 *Let k, m, and n be positive integers with $k \leq m,n$. Let (A',B') be a k-by-k bimatrix game with entries (a_{ij}, b_{ij}). Let $((x_1,\ldots,x_k),(y_1,\ldots,y_k)^t)$ be a Nash equilibrium with corresponding payoff (λ_1,λ_2). Consider (A',B') embedded as the upper left corner in an m-by-n bimatrix game (A,B) with the remaining entries (a_{ij}, b_{ij}) independently chosen with respect to some bivariate distribution with both marginals the Cauchy measure c. Let $x = (x_1,\ldots,x_k,0,\ldots,0)$ and $y = (y_1,\ldots,y_k,0,\ldots,0)^t$. Then the chance that (x,y) is an equilibrium for (A,B) is $F_c(\lambda_1)^{m-k}F_c(\lambda_2)^{n-k}$.*

Proof. For all extensions (A,B), one has simply $\lambda_1 = xAy$ and $\lambda_2 = xBy$. Let $x(i) \in \Delta^r_m$ be the i^{th} vertex so that $x(i)_u = \delta_{iu}$. Similarly let $y(j) \in \Delta^c_n$ be the j^{th} vertex so that $y(j)_u = \delta_{ju}$. Then, by considerations of convex combinations, (10) and (11) hold if and only if the extreme special cases

$$x(i)Ay \leq \lambda_1 \text{ for all } i = k+1,\ldots,m, \tag{12}$$

$$xBy(j) \leq \lambda_2 \text{ for all } j = k+1,\ldots,n \tag{13}$$

hold. For $i = k+1, \ldots, m$ and $j = 1, \ldots, k$, the quantity $a_{ij}y_j$ is distributed according to c_{y_j}. So, for each $i = k+1, \ldots, m$,

$$x(i)Ay = \sum_{j=1}^{k} a_{ij}y_j$$

is distributed according to

$$c_{y_1} * \cdots * c_{y_k} = c_{y_1+\cdots+y_k} = c_1 = c.$$

For each such i, the chance that $x(i)Ay \leq \lambda_1$ is $F_c(\lambda_1)$. Similarly, for each $j = k+1, \ldots, n$,

$$xBy(j) = \sum_{i=1}^{k} x_i b_{ij}$$

is also distributed according to c. So for each such j, the chance that $xBy(j) \leq \lambda_2$ is $F_c(\lambda_2)$. The $(m-k)$ conditions on A in (12) and the $(n-k)$ conditions on B in (13) are all independent because they involve different a_{ij} and b_{ij}. This yields the product formula given by the lemma. □

The Gaussian analog. Let g_w be the Gaussian measure with mean 0 and standard deviation w. Then one has the familiar fact $g_{w_1} * g_{w_2} = g_{\sqrt{w_1^2+w_2^2}}$. From this fact, one can deduce a lemma analogous to Lemma 1 where Cauchy is changed to Gaussian and the payoffs λ_i are replaced by the renormalized payoffs Λ_i introduced before (7). This analogous lemma plays an important role in the proof of (7).

4 Proofs of the two theorems

Proof of Theorem 1. Let $f^{c,-}_{m,n,k}(\lambda)$ be the density of Nash equilibria of support size k and payoff $(\lambda, -\lambda)$ in Cauchy-random m-by-n zero-sum games. We need to prove

$$f^{c,-}_{m,n,k}(\lambda) = \binom{m}{k}\binom{n}{k} kF_c(\lambda)^{m-1}F_c(-\lambda)^{n-1}f_c(\lambda) \tag{14}$$

for all (m,n,k) with $1 \leq k \leq m,n$. Lemma 1 says that if (14) holds for (k,k,k) then it holds for (m,n,k) for all $m,n \geq k$. By induction, we can assume that (14) holds for all (m,n,k) with k less than a given ℓ and need only prove the instance

$$f^{c,-}_{\ell,\ell,\ell}(\lambda) = \ell F_c(\lambda)^{\ell-1}F_c(-\lambda)^{\ell-1}f_c(\lambda) \tag{15}$$

of (14). We drop the superscripts c and $-$ in the rest of this proof, since they never change.

To establish (15), we prove that the $F_c(\lambda)^a$-moments of both sides coincide, with $a = 0, 1, 2, \ldots$. Taking moments of the left side of (15), and using Lemma 1 at the first step, we have

$$\begin{aligned}
&\int_{-\infty}^{\infty} f_{\ell,\ell,\ell}(\lambda) F_c(\lambda)^a d\lambda \\
&= \binom{\ell+a}{\ell}^{-1} \int_{-\infty}^{\infty} f_{\ell+a,\ell,\ell}(\lambda) d\lambda \\
&= \binom{\ell+a}{\ell}^{-1} E_{\ell+a,\ell,\ell} \\
&= \binom{\ell+a}{\ell}^{-1} \left(1 - \sum_{k=1}^{\ell-1} E_{\ell+a,\ell,k}\right) \\
&= \binom{\ell+a}{\ell}^{-1} \left(1 - \sum_{k=1}^{\ell-1} \binom{\ell+a}{k}\binom{\ell}{k} kB(\ell+a,\ell)\right). \qquad (16)
\end{aligned}$$

Taking moments of the right side of (15), via the substitution $x = F_c(\lambda)$ to obtain the standard Beta integral, we have

$$\begin{aligned}
\int_{-\infty}^{\infty} \ell F_c(\lambda)^{a+\ell-1} F_c(-\lambda)^{\ell-1} f_c(\lambda) d\lambda &= \ell \int_0^1 x^{a+\ell-1}(1-x)^{\ell-1} dx \\
&= \ell B(a+\ell, \ell). \qquad (17)
\end{aligned}$$

Setting the right sides of (16) and (17) equal, rearranging, and abbreviating $\ell + a$ by w, we see that all the $F_c(\lambda)^a$-moments of the two sides of (15) agree if and only if

$$\sum_{k=1}^{\ell} \binom{w}{k}\binom{\ell}{k} kB(w,\ell) = 1$$

holds for all $w \geq \ell$. This equation indeed holds because the k^{th} term of the left side can be interpreted as the probability that a randomly chosen ℓ-element subset of $\{1, \ldots, w+\ell-1\}$ and a randomly chosen w-element subset of $\{1, \ldots, w+\ell-1\}$ intersect in k elements. □

Proof of Theorem 2. Let $f^{c,+}_{m,n,k}(\lambda)$ be the density of Nash equilibria of support size k and payoffs (λ, λ) in Cauchy-random m-by-n coordination games. We need to prove

$$f^{c,+}_{m,n,k}(\lambda) = \binom{m}{k}\binom{n}{k} kF_c(\lambda)^{m+n-k-1} F_c(-\lambda)^{k-1} f_c(\lambda) \qquad (18)$$

for all (m, n, k). Again by Lemma 1, it suffices to prove the special case of equal indices,

$$f^{c,+}_{\ell,\ell,\ell}(\lambda) = \ell F_c(\lambda)^{\ell-1} F_c(-\lambda)^{\ell-1} f_c(\lambda). \qquad (19)$$

Comparing (19) with (15), we see that it suffices to prove

$$f^{c,+}_{\ell,\ell,\ell}(\lambda) = f^{c,-}_{\ell,\ell,\ell}(\lambda) \qquad (20)$$

for all ℓ.

Now in general, suppose (A, B) is a ℓ-by-ℓ bimatrix game with only finitely many Nash equilibria. Let A^{adj} and B^{adj} be the corresponding adjoint matrices, so that $A^{\mathrm{adj}} = \det(A)A^{-1}$ and $B^{\mathrm{adj}} = \det(B)B^{-1}$ when the right sides are defined. A necessary condition for (A, B) to have a Nash equilibrium of support size ℓ is that the row sums of A^{adj} all have the same sign and the column sums of B^{adj} all have the same sign. Then (A, B) has a unique Nash equilibrium of support size ℓ, given by the formulas

$$x = r_\ell B^{\mathrm{nadj}}, \qquad y = A^{\mathrm{nadj}} c_\ell.$$

Here B^{nadj} is the normalized adjoint of B, i.e. $B^{\mathrm{adj}}/\sigma(B^{\mathrm{adj}})$ where $\sigma(B^{\mathrm{adj}})$ is the sum of the entries of B^{adj}; similarly, $A^{\mathrm{nadj}} = A^{\mathrm{adj}}/\sigma(A^{\mathrm{adj}})$. Also r_ℓ and c_ℓ are row and column ℓ-vectors with all entries 1. With these definitions, x_i is the i^{th} column sum of B^{nadj} and y_j is the j^{th} row sum of A^{nadj}.

From the previous paragraph, one sees that (A, B) has a unique Nash equilibrium with support size ℓ if and only if $(A, -B)$ does. In this case, the equilibrium (x, y) for (A, B) is also an equilibrium for $(A, -B)$, with the payoff to Player 1 being xAy in both cases. Taking $B = A$ relates the zero-sum case to the coordination case in the way we need to establish (20). In fact, our argument shows that (20) holds with c replaced by any measure μ with an even density function. $\square$

The case $t = 0$. McLennan considered the chance that a Gaussian-random ℓ-by-ℓ neutral game has a Nash equilibrium with support size ℓ and normalized payoffs (Λ_1, Λ_2) in a rectangle R. He proved that this chance is the integral of

$$f^{g,0}_{\ell,\ell,\ell}(\Lambda_1, \Lambda_2) = \frac{1}{2^{2\ell-2}} \left(\sqrt{\frac{\ell}{2\pi}} e^{-\ell \Lambda_1^2/2} \right) \left(\sqrt{\frac{\ell}{2\pi}} e^{-\ell \Lambda_2^2/2} \right) \tag{21}$$

over R. From this special case, and the Gaussian lemma discussed at the end of the last section, one gets the full $(g, 0)$ case underlying (7). The obstruction to likewise establishing the $(c, 0)$ case is that there does not seem to be a formula for $f^{c,0}_{\ell,\ell,\ell}(\lambda_1, \lambda_2)$ analogous to (21).

5 Two corollaries describing asymptotics

In this section, we derive the asymptotic behavior associated to our two theorems in the limit of large matrices of a given shape. We work first at a basic level corresponding to means and then at a more refined level corresponding to variances. In terms of Figure 1, the basic level gives the central point of the clouds and the qualitative fact that the clouds contract on their central points as matrix sizes increase. The refined level gives more information on the shape of the clouds and how fast they contract.

Basic asymptotics: first derivatives and means. It is best to change variables and work with s, r_1, r_2, κ_1, κ_2, κ, and ν, rather than m, n, k, and λ. The *size* of an m-by-n matrix is $s = m+n$. The *shape* of an m-by-n matrix is the pair $(r_1, r_2) = (m, n)/s$. If a Nash equilibrium of an m-by-n bimatrix game has support size k, then the *support fractions* of Player 1 and 2 are respectively $\kappa_1 = k/m$ and $\kappa_2 = k/n$. The *total support fraction* is $\kappa = \kappa_1 + \kappa_2$. Finally, $\nu = 2/\pi \arctan(\lambda)$ is a convenient renormalization of the payoff λ.

To pass from the old variables to the new ones, we use the formulas,

$$m = r_1 s, \quad n = r_2 s, \quad k = r_1 r_2 \kappa s = r_1 \kappa_1 s = r_2 \kappa_2 s, \quad \lambda = \tan(\pi\nu/2).$$

There are three relations among our seven new variables,

$$r_1 + r_2 = 1, \qquad \kappa_1 = r_2 \kappa, \qquad \kappa_2 = r_1 \kappa.$$

We view s, r_1, κ, and ν as our main new variables. However we systematically use r_2, κ_1, and κ_2 as abbrevations, so as to keep symmetries and other intuitive underpinnings of formulas as evident as possible.

Let R be the rectangle

$$0 < \kappa \leq \min(\frac{1}{r_1}, \frac{1}{r_2}), \qquad -1 < \nu < 1 \tag{22}$$

in the κ-ν plane. A Nash equilibrium of an m-by-n game gives rise to a point (κ, ν) in R. The corollaries we are presently deriving describe the distribution of these points as s tends to ∞.

Write $(e_1, e_2) = (r_1, r_2)$ in the zero-sum case $t = -$ and $(e_1, e_2) = (1 - r_1 r_2 \kappa, r_1 r_2 \kappa)$ in the coordination case $t = +$. Translating the formulas of Theorems 1 and 2 into the present notation gives measures μ_s^t on R defined by

$$\mu_s^t\left([\kappa', \kappa''] \times [\nu', \nu'']\right) = \sum_{k=\lceil r_1 r_2 s \kappa' \rceil}^{\lfloor r_1 r_2 s \kappa'' \rfloor} \binom{r_1 s}{k} \binom{r_2 s}{k} k \int_{\nu'}^{\nu''} \left(\frac{1+\nu}{2}\right)^{e_1 s - 1} \left(\frac{1-\nu}{2}\right)^{e_2 s - 1} d\nu. \tag{23}$$

Here the summation is over integers satisfying $s r_1 r_2 \kappa' \leq k \leq s r_1 r_2 \kappa''$. Each measure μ_s^t has support the union of the vertical lines $\kappa = k/(r_1 r_2 s)$ for $k = 1, \ldots, \min(r_1 s, r_2 s)$. Figure 1 in Section 1 can be viewed as a window on R. The plotted dots roughly indicate μ_{200}^- and μ_{200}^+ in the case $r_1 = r_2 = 1/2$. Vertical lines, spaced $1/(r_1 r_2 s) = 0.02$ apart, are visually evident.

Using Stirling's asymptotic formula $s! \sim s^{s+1/2} e^{-s} \sqrt{2\pi}$ to eliminate factorials, one formally gets an approximation for large s,

$$\mu_s^t \approx s K(\kappa, \nu) M^t(\kappa, \nu)^s d\kappa\, d\nu. \tag{24}$$

Here

$$K(\kappa, \nu) = \frac{r_1 r_2}{\pi \sqrt{(1-\kappa_1)(1-\kappa_2)}} \cdot \frac{1}{(1-\nu)(1+\nu)}, \tag{25}$$

$$M^t(\kappa, \nu) = \frac{(1-\kappa_1)^{r_1(\kappa_1 - 1)} (1-\kappa_2)^{r_2(\kappa_2 - 1)}}{\kappa_1^{r_1 \kappa_1} \kappa_2^{r_2 \kappa_2}} \cdot \frac{(1+\nu)^{e_1} (1-\nu)^{e_2}}{2}. \tag{26}$$

Equation (34) below rewrites (24) in a slightly stronger form and there we explain the rigorous sense in which "≈" is to be understood. For the moment, the formal statement (24) makes clear that it is important to find the global maxima of $M^t(\kappa,\nu)$ on R. To do this, we carry out a critical point analysis.

The first partials of the quantities $M^t(\kappa,\nu)$ are most easily obtained via logarithmic differentiation and work out as follows.

$$M_\kappa^-(\kappa,\nu) = r_1 r_2 M^-(\kappa,\nu)\log\left(1+\frac{1-\kappa}{r_1 r_2 \kappa^2}\right), \tag{27}$$

$$M_\nu^-(\kappa,\nu) = \frac{M^-(\kappa,\nu)}{(1-\nu)(1+\nu)}\left(r_1 - r_2 - \nu\right), \tag{28}$$

$$M_\kappa^+(\kappa,\nu) = r_1 r_2 M^+(\kappa,\nu)\log\left(1+\frac{(1-\kappa)+(-1+\kappa-2r_1r_2\kappa^2)\nu}{r_1 r_2\kappa^2(1+\nu)}\right), \tag{29}$$

$$M_\nu^+(\kappa,\nu) = \frac{M^+(\kappa,\nu)}{(1+\nu)(1-\nu)}\left(1-\nu-2\kappa r_1 r_2\right). \tag{30}$$

In each of the four cases, the partial vanishes if and only if its last factor does. It is clear from (27) and (28) that $M^-(\kappa,\nu)$ has a unique critical point, namely

$$\langle\kappa\rangle^{c,-} = 1, \qquad \langle\nu\rangle^{c,-} = r_1 - r_2. \tag{31}$$

For the case $t = +$, algebra shows that again there is a unique critical point. It is

$$\langle\kappa\rangle^{c,+} = \frac{1}{1+Q}, \qquad \langle\nu\rangle^{c,+} = Q = \sqrt{1-2r_1r_2}. \tag{32}$$

These four quantities play a central role in the rest of this paper. They appear as limiting means in Corollaries 1 and 2. They are graphed together with analogs in Figures 5 and 6.

Substituting the critical point (31) into the function (26), one gets a complicated expression which simplifies all the way to give the critical value $V^{c,-} = 1$ in the zero-sum case. Substituting (32) into (26), one gets that the the critical value in the coordination case is

$$V^{c,+} = 2^{-1/2} r_2^{r_2-r_1}(1+Q)(1-3r_1r_2-(r_1-r_2)Q)^{(r_1-r_2)/2}. \tag{33}$$

Figure 2 visually summarizes our critical point analysis so far, in the square case $(r_1,r_2) = (1/2,1/2)$. To facilitate comparison with Figure 1, we use the variable $\kappa_1 = \kappa/2$. The critical point is at $(0,0)$ and $(1-2^{-1/2}, 2^{-1/2}) \approx (0.29, 0.71)$ in the (κ_1,ν) plane, in the zero-sum and coordination case respectively. The critical values are 1 and $1/2 + 1/\sqrt{2} \approx 1.2071$ respectively. One can rigorously check that the critical value found is indeed the unique maximum by comparing values with limiting values on the boundary of R.

Refined asymptotics: second derivatives and variances. We now drop the superscript c, to lighten notation. We drop the superscripts $-$ and $+$ as

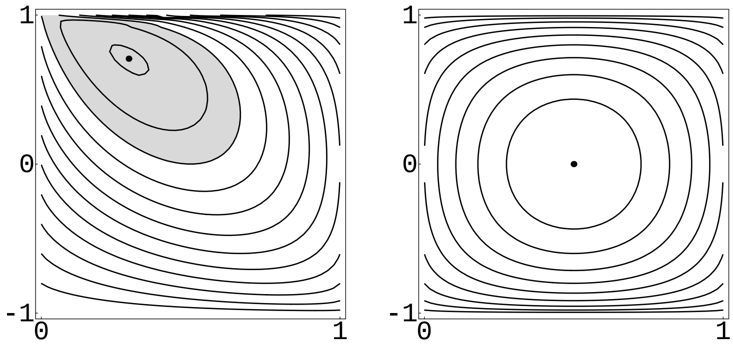

Fig. 2 Contour graphs of $M^{+}(2\kappa_1,\nu)$ and $M^{-}(2\kappa_1,\nu)$ in the square case $r_1 = r_2 = 1/2$, with contours spaced by 0.1. The gray region is where $M^{+}(2\kappa_1,\nu) \geq 1$ and represents where a large Cauchy-random coordination game has Nash equilibria. The unique critical point is indicated in each case.

well, as the distinction will be clear. We indicate critical values $f(\langle\kappa\rangle,\langle\nu\rangle)$ simply by f^{crit}.

Our word "refined" refers to the fact that we "zoom in" on the κ-ν plane at the critical point $(\langle\kappa\rangle,\langle\nu\rangle)$ by using the change of variables $x = \sqrt{s}(\kappa - \langle\kappa\rangle)$ and $y = \sqrt{s}(\nu - \langle\nu\rangle)$. Consider now four families of measures on the new x-y plane, each dependent on s as follows.

$$\mu_s \approx K\left(\langle\kappa\rangle + \frac{x}{\sqrt{s}}, \langle\nu\rangle + \frac{y}{\sqrt{s}}\right) M\left(\langle\kappa\rangle + \frac{x}{\sqrt{s}}, \langle\nu\rangle + \frac{y}{\sqrt{s}}\right)^{s} dx\,dy, \quad (34)$$

$$\approx K^{\mathrm{crit}}\left(V + \frac{M^{\mathrm{crit}}_{\kappa\kappa}x^2 + 2M^{\mathrm{crit}}_{\kappa\nu}xy + M^{\mathrm{crit}}_{\nu\nu}y^2}{2s}\right)^{s} dx\,dy, \quad (35)$$

$$\approx K^{\mathrm{crit}}V^{s}\exp\left(\frac{M^{\mathrm{crit}}_{\kappa\kappa}x^2 + 2M^{\mathrm{crit}}_{\kappa\nu}xy + M^{\mathrm{crit}}_{\nu\nu}y^2}{2V}\right) dx\,dy. \quad (36)$$

Here the first approximation is a rewriting of (24), shifted from the κ-ν plane to the x-y plane. An s has dropped out because of the relation $dx\,dy = s\,d\kappa\,d\nu$. Each of the three approximations holds in the sense that, applied to any rectangle $[x', x''] \times [y', y'']$, the limiting ratio of the left and right sides tends to 1 as $s \to \infty$. Approximations (34), (35), and (36) hold respectively because of the sufficient accuracy of Stirling's approximation, quadratic approximation, and the asymptotic formula $\exp(z) \sim (1 + z/s)^{s}$.

To proceed further, we need to calculate the second critical partials appearing in (36). The right sides of (27)-(30) are all written in the form $A(\kappa,\nu)B(\kappa,\nu)$ with $B(\kappa,\nu)$ the last printed factor. As already mentioned, $B(\langle\kappa\rangle,\langle\nu\rangle) = 0$ in each case. Thus critical second derivatives can always be calculated by $A(\langle\kappa\rangle,\langle\nu\rangle)B'(\langle\kappa\rangle,\langle\nu\rangle)$, as the other term in the product formula vanishes.

In the zero-sum case, one finds

$$K^{\text{crit}} = \frac{1}{4\pi\sqrt{r_1 r_2}}, \qquad \left(M_{\kappa\kappa}^{\text{crit}},\ M_{\kappa\nu}^{\text{crit}},\ M_{\nu\nu}^{\text{crit}}\right) = \left(-1,\ 0,\ \frac{-1}{4r_1r_2}\right).$$

To state our corollary in terms of the payoff λ rather than the adjusted payoff ν, we use $\lambda = \tan(\pi\nu/2)$ and convert via $F = (d\lambda/d\nu)^{\text{crit}}$.

Corollary 1 *Consider Cauchy-random zero-sum games with increasing size $s = m+n$ and fixed shape $(r_1, r_2) = (m, n)/s$. The probability that the unique Nash equilibrium has support fractions $(\kappa_1, \kappa_2) = (k/m, k/n) = (r_2, r_1)\kappa$ and payoffs $(\lambda, -\lambda)$ satisfying*

$$\frac{a}{\sqrt{s}} \leq \kappa - 1 \leq \frac{b}{\sqrt{s}} \quad \textit{and} \quad \frac{c}{\sqrt{s}} \leq \lambda - \tan\left(\frac{\pi}{2}(r_1 - r_2)\right) \leq \frac{d}{\sqrt{s}}$$

is asymptotic to

$$\frac{F}{4\pi\sqrt{r_1r_2}} \int_a^b \int_c^d \exp\left(-\frac{x^2}{2} - \frac{F^2y^2}{8r_1r_2}\right) dx\, dy. \tag{37}$$

where $F = \frac{\pi}{2}\sec^2\left(\frac{\pi}{2}(r_1 - r_2)\right)$.

Verbally, the corollary says that the rescaled quantities $x = \sqrt{s}(\kappa - \langle\kappa\rangle)$ and $Fy = \sqrt{s}(\lambda - \langle\lambda\rangle)$ are both normally distributed in the limit of large s. As an immediate but much cruder consequence, the total support fraction κ is asymptotically distributed according to the point mass at $\langle\kappa\rangle$ while the payoff λ is asymptotically distributed according to the point mass at $\langle\lambda\rangle$.

Before proceeding to the coordination case, note that

$$\int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \exp\left(-(x_1, \ldots, x_d)P(x_1, \ldots, x_d)^t\right) dx_1 \cdots dx_d = \frac{\pi^{d/2}}{\sqrt{\det(P)}} \tag{38}$$

for P a positive definite symmetric d-by-d matrix. This Gaussian integral formula with $d = 2$ lets one check by inspection that (37) evaluates to 1 over the whole plane, as it must. It also is the tool for passing from (39) to (40) below.

In the coordination case, the quantities needed in (36) work out to

$$K^{\text{crit}} = \frac{1}{\pi\sqrt{2}}, \quad \left(M_{\kappa\kappa}^{\text{crit}},\ M_{\kappa\nu}^{\text{crit}},\ M_{\nu\nu}^{\text{crit}}\right) = \left(-2r_1r_2(1+2Q),\ -1,\ \frac{-1}{2r_1r_2}\right)V.$$

To state our corollary in terms of the payoff λ rather than the adjusted payoff ν, we again convert via $F = (d\lambda/d\nu)^{\text{crit}}$.

Corollary 2 *Consider Cauchy-random coordination games with increasing size $s = m + n$ and fixed shape $(r_1, r_2) = (m, n)/s$. The expected number* *of Nash equilibria with support fractions* $(\kappa_1, \kappa_2) = (k/m, k/n) = (r_2, r_1)\kappa$ *and payoffs* (λ, λ) *satisfying*

$$\frac{a}{\sqrt{s}} \le \kappa - \frac{1}{1+Q} \le \frac{b}{\sqrt{s}} \quad \textit{and} \quad \frac{c}{\sqrt{s}} \le \lambda - \tan(\frac{\pi}{2}Q) \le \frac{d}{\sqrt{s}}$$

is asymptotic to

$$\frac{FV^s}{\sqrt{2}\pi} \int_a^b \int_c^d \exp\left(-r_1 r_2(1+2Q)x^2 - Fxy - \frac{F^2y^2}{4r_1r_2},\right) dxdy \tag{39}$$

where $F = \frac{\pi}{2}\sec^2\left(\frac{\pi}{2}Q\right)$. *In particular, the mean total number of Nash equilibria is*

$$\sum_{k=1}^{\min(m,n)} E^{c,+}_{m,n,k} = \frac{1}{\sqrt{Q}} V^s. \tag{40}$$

Here Q *and* V *depend on the shape* (r_1, r_2) *via (32b), and (33) respectively.*

In rough summary, Cauchy-random r_1s-by-r_2s coordination games have a very large number of Nash equilibria if s is large. The vast majority of these equilibria behave similarly in that their (κ, λ) are very close to $(\langle\kappa\rangle, \langle\lambda\rangle)$. As a numerical example of (40), Theorem 2 says that 100-by-100 Cauchy-random coordination games have approximately 2.659×10^{16} Nash equilbria, and Corollary 2 approximates this with $V^{200}/\sqrt{Q} \approx 2.657 \times 10^{16}$.

Asymptotics for the McLennan-Berg theorem. We conclude this section by applying our techniques to the McLennan-Berg theorem to obtain sharper statements than those given in McLennan and Berg (2005). From the discussion below, the only parts explicitly appearing in this reference are the numerical approximations on the right sides of (44a) and (44b).

Analogously to (24), Nash equilibria have a density asymptotic to a function $sK(\kappa)M^0(\kappa, \Lambda_1, \Lambda_2)^s$. In contrast to (5) and (8), the exponents under the integrals of (7) depend linearly on s. It is for this reason that $K(\kappa)$ depends only on κ and not on Λ_1 or Λ_2. In fact, $K(\kappa)$ is $2^{3/2}/\pi$ times the first fraction of $K(\kappa, \nu)$ appearing in (25).

The quantity $M^0(\kappa, \Lambda_1, \Lambda_2)$ is likewise quite similar to the previous $M^t(\kappa, \nu)$. From the fact that there is a product of two one-variable integrals in (7), one gets a factorization $M^0(\kappa, \Lambda_1, \Lambda_2) = M^0(\kappa_1, \Lambda_1)M^0(\kappa_2, \Lambda_2)$. The factors work out to

$$M^0(\kappa_i, \Lambda_i) = \frac{(1-\kappa_i)^{r_i(\kappa_i-1)}}{\sqrt{2}\kappa_i^{r_i\kappa_i}} \left(1 + \operatorname{erf}\left(\frac{\Lambda_i}{\sqrt{2}}\right)\right)^{r_i(1-\kappa_i)} e^{-r_i\kappa_i\Lambda_i^2/2}. \tag{41}$$

Comparing with (26), one sees that the first fraction there appears also in $M^0(\kappa, \Lambda_1, \Lambda_2)$.

As in our Cauchy cases, the function to be maximized depends on the parameter r_1. However the situation here is more complicated as the current function $M^0(\kappa, \Lambda_1, \Lambda_2)$ depends on three variables, rather than the two

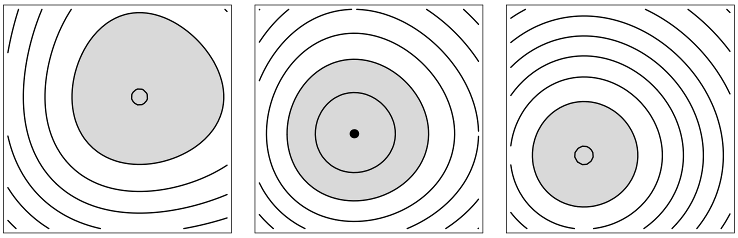

Fig. 3 Contour graphs of $M^0(2\kappa_1, \Lambda_1, \Lambda_2)$ corresponding to the square case $r_1 = r_2 = 1/2$ and the window $\Lambda_1, \Lambda_2 \in [-0.5, 2.5]$ on the Λ_1-Λ_2 plane. From left to right, the three fixed values of κ_1 are 0.127, $\langle\kappa_1\rangle \approx 0.316$, and 0.503. Following previous conventions, contours are spaced by 0.1 and the gray region is where $M^0(2\kappa_1, \Lambda_1, \Lambda_2) \geq 1$. The maximum value in the center figure is the critical value $V \approx 1.1512$. The left and right values of κ_1 are chosen such that the maximum value of $M^0(2\kappa_1, \Lambda_1, \Lambda_2)$ is just slightly over 1.1.

variables κ and ν. Figure 3 is as analogous as possible to Figure 2, given this difference. For general r_1, the unique critical point $(\langle\kappa\rangle, \langle\Lambda_1\rangle, \langle\Lambda_2\rangle)$ is a global maximum and can be found as the solution to

$$\Lambda_1\Lambda_2 = 2/\pi, \qquad \kappa = \kappa_1 + \kappa_2, \tag{42}$$

$$\kappa_i = \frac{e^{-\Lambda_i^2/2}}{1 + \sqrt{\pi/2}\Lambda_i\left(1 + \mathrm{erf}(\Lambda_i/\sqrt{2})\right)}, \qquad r_i = 1 - \frac{\kappa_i}{\kappa}. \tag{43}$$

These equations are best obtained from $M^0(\kappa, \Lambda_1, \Lambda_2)$ by logarithmic differentiation, as in the Cauchy case.

In the square case $(r_1, r_2) = (1/2, 1/2)$, solving the system (42), (43) gives $\langle\Lambda_1\rangle = \langle\Lambda_2\rangle = \sqrt{2/\pi}$. Abbreviating $q = e^{1/\pi}(1+\mathrm{erf}(1/\sqrt{\pi})) \approx 2.1654$, one has further

$$\langle\kappa_1\rangle = \frac{1}{1+q} \approx 0.3159 \qquad V = \frac{1+q}{2e^{1/\pi}} \approx 1.1512. \tag{44}$$

In non-square cases, it does not seem possible to solve (42), (43) for the three variables κ, Λ_1, and Λ_2 to get expressions for $\langle\kappa\rangle$, $\langle\Lambda_1\rangle$, and $\langle\Lambda_2\rangle$ as explicit classical functions of r_1. However one can work with $\langle\Lambda_1\rangle \in (0, \infty)$ as a parameter, and then very easily express r_1, $\langle\kappa\rangle$, $\langle\Lambda_2\rangle$ in terms of it. This suffices for drawing the corresponding dashed curves in Figures 4 and 5.

As in the Cauchy cases, computations with second derivatives show that the standard deviations of κ, Λ_1, Λ_2 about their respective limiting means $\langle\kappa\rangle$, $\langle\Lambda_1\rangle$, $\langle\Lambda_2\rangle$ decay as $1/\sqrt{s}$. To obtain a formula analogous to (40), we use the three-variable analog of (36). There is an extra factor of $\sqrt{s}$ in the denominator of this analog because of the extra variable. One gets that the mean total number of Nash equilibria has the form

$$\sum_{k=1}^{\min(m,n)} E^{g,0}_{m,n,k} \sim \frac{C}{\sqrt{s}} V^s. \tag{45}$$

Using (38) in $d = 3$ dimensions, one gets $C = K^{\text{crit}}\pi^{3/2}/\sqrt{\det(M^{\text{crit}}/V)}$ where M^{crit} is the symmetric three-by-three matrix of critical second partials of $M^0(\kappa, \Lambda_1, \Lambda_2)$. In the square case $(r_1, r_2) = (1/2, 1/2)$, one gets

$$C = \frac{2\sqrt{2}(1+q)}{\sqrt{2+2q+\pi q}} \approx 2.4705. \tag{46}$$

As a numerical example, the McLennan-Berg theorem says that 100-by-100 games have on average about 2.982×10^{11} Nash equilibria while (45) gives the approximation 2.978×10^{11}.

6 Conclusion

We conclude by graphically comparing the main quantities arising in the asymptotic analyses of the previous section. We explain in informal game-theoretic terms how the three cases fit into a coherent whole. As in the previous section, we consider m-by-n matrices of large size $s = m + n$ and fixed shape $(r_1, r_2) = (m, n)/s$. Our discussion centers on three figures, each graphing functions of r_1. Quantities based on the Cauchy and Gaussian measures are graphed with solid and dashed curves respectively.

Total number of Nash equilibria. The total number of Nash equilibria in the cases $(c, +)$, and $(g, 0)$ is asymptotic to $C^{c,+}(V^{c,+})^s$ and $C^{g,0}(V^{g,0})^s/\sqrt{s}$ by (40) and (45) respectively.

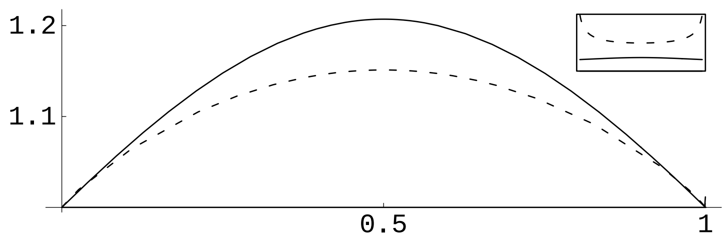

Fig. 4 Functions giving the total number of Nash equilibria. From bottom to top, $V^{\mu,-} = 1$, $V^{g,0}$, and $V^{c,+}$. In the inset, $C^{c,+}$ on the bottom and $C^{g,0}$ on the top.

The main part of Figure 4 graphs $V^{c,+}$ and $V^{g,0}$. The horizontal axis can be understood as $V^{\mu,-}$, corresponding to the fact that a μ-random m-by-n zero-sum game has exactly one Nash equilibrium, independent of μ, m, and n. The inset of Figure 4 graphs the functions $C^{c,+}$ and $C^{g,0}$ with vertical range $[0, 5]$. The function $C^{c,+}$ has values ranging from 1 at the endpoints to $2^{1/4} \approx 1.189$ in the middle, while $C^{g,0}$ approaches ∞ at each endpoint.

The clear import of Figure 4 is that for fixed shape (r_1, r_2) and sufficiently large size s, as one passes from the extreme of zero-sum games to the extreme of coordination games, the mean number of Nash equilibria rapidly increases. This is a strong phenomenon for all shapes (r_1, r_2), and it is strongest in the square case $(r_1, r_2) = (1/2, 1/2)$.

Support fractions. In each of the three situations $(c, -)$, $(g, 0)$, and $(c, +)$, Player 1's support fractions $\kappa_1 = k/m$ cluster closely about the corresponding limiting mean. Figure 5 graphs the three limiting means.

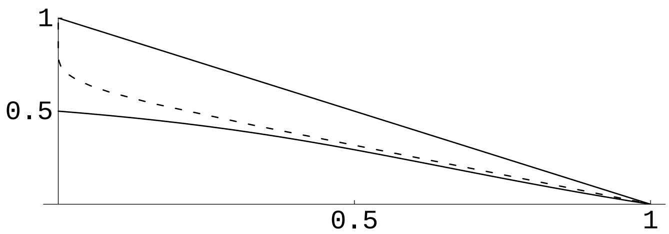

Fig. 5 Limiting mean support fractions for Player 1. From top to bottom, $\langle \kappa_1 \rangle^{c,-} = 1 - r_1$, $\langle \kappa_1 \rangle^{g,0}$, and $\langle \kappa_1 \rangle^{c,+}$.

Berg and Engel (1998) have investigated the case $(g, -)$. With the help of [3], we have numerically solved their Equations 6. From these computations, it seems that always $0.969 < \langle \kappa_1 \rangle^{g,-} / \langle \kappa_1 \rangle^{c,-} \leq 1$, with equality both in the square case $r_1 = 1/2$ and, in a limiting sense, at the endpoints $r_1 = 0$ and $r_1 = 1$. This comparison gives one confidence that in some situations the exact definition of randomness is of only secondary importance. Numerical computation with small games in the settings $(c, 0)$ and $(g, +)$ increases confidence further.

The simple qualitative import of Figure 5 is as follows. When r_1 is near zero in a zero-sum game, Player 1 is disadvantaged in a competitive environment. His best defense is to suitably mix nearly all his pure strategies. As either Player 1 becomes less disadvantaged or as the environment becomes less competitive, Player 1 plays a smaller fraction of his pure strategies in Nash equilibria.

Payoffs. For payoffs, we consider the Cauchy case only. The Gaussian case is qualitatively similar, but quantitative comparison across measures requires that one enter into a number of scaling issues. In each of the two situations $(c, -)$ and $(c, +)$, Player 1's payoffs λ cluster closely about the corresponding limiting mean. Figure 6 graphs the limiting means $\langle \lambda \rangle^{c,-}$ and $\langle \lambda \rangle^{c,+}$, and also an approximation to $\langle \lambda \rangle^{c,0}$, obtained by extrapolating from games with $s \leq 20$.

Figure 6 confirms intuitive expectations. For any given r_1, as the correlation between the interests of the players increases, the return to each player at an average Nash equilibrium increases. In the zero-sum and neutral cases, as r_1 increases, the return to Player 1 increases as well. The fact that $\langle \lambda \rangle^{c,+}$ has a global minimum at $r_1 = 1/2$ corresponds to the fact that

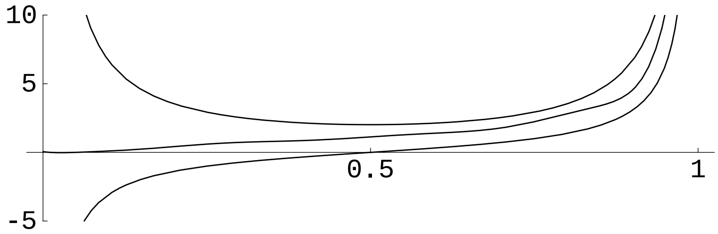

Fig. 6 Limiting mean payoffs to Player 1. From bottom to top, $\langle\lambda\rangle^{c,-} = -\cot(\pi r_1)$, an approximation to $\langle\lambda\rangle^{c,0}$, and $\langle\lambda\rangle^{c,+}$.

coordination when cooperation is disallowed is difficult: it is better for both players if one player has most of the control.

References

1. Berg J (2000), Statistical mechanics of random two-player games, Phys. Review E (3) 61: 2327–2339.
2. Berg J, Engel A (1998) Matrix games, mixed strategies, and statistical mechanics. Physics Review Letters 81: 4999–5002.
3. Berg J, Engel A (2004) e-mail communications to the author
4. Berg J, Weigt M (1999) Entropy and typical properties of Nash equilibria in two-player games. Europhysics Letters 48 (2), 129-135.
5. Dwass, M (1985) On the convolution of Cauchy distributions, Amer. Math. Monthly 92 no. 1: 55–57.
6. Faris WG, Maier RS (1987) The value of a random game: the advantage of rationality. Complex Systems 1 no. 2: 235–244.
7. Goldman AJ (1957) The probability of a saddle point. American Math Monthly 64: 729–730.
8. McLennan A (2005) The expected number of Nash equilibria of a normal form game. Econometrica 73, no. 1: 141–174.
9. McLennan A, Berg J (2005) The asymptotic expected number of Nash equilibria of two-player normal form games. Games Econ. Behav. 51, no. 2: 264–295.
10. Nelson, RB (1985) Letter to the editor, Amer. Math. Monthly 92 no. 9: 679.
11. Powers I (1990) Limiting distributions on the number of pure strategy Nash equilibria. Internat. J. Game Theory 19: 277–286.
12. Roberts DP (2005) Pure Nash equilibria of coordination games. Econom. Lett. 89, no. 1: 7–11.
13. Rose C, Smith MD (2002) Mathematical Statistics with *Mathematica*, Springer Texts in Statistics.
14. Stanford W (1995) A note on the probability of k pure Nash equilibria in matrix games. Games Econ. Behav. 9: 238–246.
15. Stanford W (1999) On the number of pure strategy Nash equilibria in finite common payoffs games, Econom. Lett. 62: 29–34.
16. Thrall RM, Falk JE (1965) Some results concerning the kernel of a game. SIAM Review, 7: 359–375.