

Motor coordination: when two have to act as one

Daniel A. Braun · Pedro A. Ortega ·
Daniel M. Wolpert

Received: 18 November 2010 / Accepted: 15 March 2011 / Published online: 1 April 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract Trying to pass someone walking toward you in a narrow corridor is a familiar example of a two-person motor game that requires coordination. In this study, we investigate coordination in sensorimotor tasks that correspond to classic coordination games with multiple Nash equilibria, such as “choosing sides,” “stag hunt,” “chicken,” and “battle of sexes”. In these tasks, subjects made reaching movements reflecting their continuously evolving “decisions” while they received a continuous payoff in the form of a resistive force counteracting their movements. Successful coordination required two subjects to “choose” the same Nash equilibrium in this force-payoff landscape within a single reach. We found that on the majority of trials coordination was achieved. Compared to the proportion of trials in which miscoordination occurred, successful coordination was characterized by several distinct features: an increased mutual information between the players’ movement endpoints, an increased joint entropy during the movements, and by differences in the timing of the players’ responses. Moreover, we found that the probability of successful coordination depends on the players’ initial distance from the Nash equilibria. Our results suggest that two-person coordination arises naturally in motor interactions and is facilitated by favorable initial positions, stereotypical motor pattern, and differences in response times.

Keywords Motor coordination · Multi-agent interaction · Game theory

Introduction

Many human interactions require coordination. When we walk down a corridor, for example, and encounter someone walking toward us, there are two equally viable options in this classic “choosing sides” game: we could move to the right while the other person passes on the left or conversely move to the left while the other person passes on the right. Crucially, each of the two solutions requires that the two actors make their choice in a mutually consistent fashion. In everyday situations, we usually master this problem gracefully, although sometimes it can lead to a series of awkward readjustments. Such continuous pair-wise interactions are distinct from the typical studies of cognitive game theory in which subjects may be required to make single discrete choices such as choosing left or right. In such a discrete situation, without communication or a previous agreement, the coordination problem can initially only occasionally be solved by chance, if both players happen to make the correct choice.

The problem of continuous motor coordination has been previously studied from a dynamical systems point of view (Kelso 1995), but not within the framework of game theory. Coordination problems have a long history in game theory and the social sciences (Fudenberg and Tirole 1991). In the eighteenth century, for example, Rousseau considered the now classic “stag hunt” game. In this game, each member of a group can choose to either hunt a highly valued stag or a lower-valued hare. However, the stag is caught only if everybody cooperates to hunt it in unison whereas each person can capture a hare independent of

Electronic supplementary material The online version of this article (doi:10.1007/s00221-011-2642-y) contains supplementary material, which is available to authorized users.

D. A. Braun (✉) · P. A. Ortega · D. M. Wolpert
Department of Engineering, Computational and Biological
Learning Laboratory, University of Cambridge, Cambridge, UK
e-mail: dab54@cam.ac.uk

cooperation (Rousseau et al. 2008). This coordination problem has two different solutions: the cooperative stag solution which is called payoff-dominant, since it gives higher payoff to all participants, and the hare solution which is risk-dominant, because it is the safe option that cannot be thwarted by anybody.

Unlike games such as prisoners' dilemma which have a single Nash solution, coordination games have multiple Nash equilibria. In a Nash equilibrium, each player chooses a strategy so that no player has anything to gain by changing only his or her strategy unilaterally. As no player has any incentive to switch their strategy, each player's strategy in a Nash equilibrium is a best response given the strategies of all the other players. Nash equilibria can be pure or mixed. In a pure Nash equilibrium, each player chooses a deterministic strategy that uniquely determines every move the player will make for any situation he or she could face. In a mixed Nash equilibrium, players stochastically choose their strategies, thus, randomizing their actions so that they cannot be predicted by their opponent. In coordination games, there are multiple Nash equilibria, often both pure and mixed. Importantly, although the Nash equilibrium concept (Nash 1950) specifies which solutions should occur in a game, in coordination games the problem of how to select a particular equilibrium—that is how to coordinate—is unanswered. To address this issue, the concept of correlated equilibrium has been proposed (Aumann 1974, 1987), in which the two players accept a coordinating advice from an impartial third party so as to correlate their decisions. Consider for example the coordination game of “battle of sexes,” where a husband enjoys taking his wife to a football match (which she dislikes) while she prefers to take her husband to the opera (which he abhors). However, both dislike going on their own even more. A correlated equilibrium would then allow a solomonic solution of taking turns as indicated by the third party. Similarly, in the game of “chicken,” which is also known as the game of “hawk and dove,” two players do not want to yield to the other player, however, the outcome where neither player yields is the worst possible for both players. This game is often illustrated by two drivers that are on collision course, where neither wants to swerve in order not to be called a “chicken,” however, if neither swerves then both suffer a great loss. Again an impartial third party could determine who swerves and who stays over repeated trials to make it a “fair” game.

Correlated equilibria can also be reached in the absence of an explicit third party by learning over many trials, provided that the same coordination problem is faced repeatedly (Chalkiadakis and Boutilier 2003; Greenwald and Hall 2003). In the case of continuous motor games, coordination might even be achievable within a single trial through the process of within-trial adaptation in continuous

motor interactions (Braun et al. 2009a). In the following, we will refer to this “within-trial” adaptation simply as adaptation reflecting changes in subjects' responses during individual trials or reaches. To investigate two-player coadaptation in motor coordination tasks, we employed a recently developed task design that translates classical 2-by-2 matrix games into continuous motor games (Braun et al. 2009b). We exposed subjects to the motor versions of four different coordination games: “choosing sides,” “stag hunt,” “chicken,” and “battle of sexes”. In each game, there were always two pure Nash equilibria and one mixed Nash equilibrium. The presence of multiple equilibria is the crucial characteristic of coordination games since it poses the problem of equilibrium selection or coordination. This allowed us to study the general features that characterize successful coordination in a number of different two-player motor interactions.

Methods

Participants

Twelve naïve participants provided written informed consent before taking part in the experiment in six pairs. All participants were students at the University of Cambridge. The experimental procedures were approved by the Psychology Research Ethics Committee of the University of Cambridge.

Experimental apparatus

The experiment was conducted using two vBOTs which are planar robotic interfaces. The custom-built device consists of a parallelogram constructed mainly from carbon fiber tubes that are driven by rare earth motors via low-friction timing belts. High-resolution incremental encoders are attached to the drive motors to permit accurate computation of the robot's position. Care was taken with the design to ensure it was capable of exerting large end-point forces while still exhibiting high stiffness, low friction, and also low inertia. The robot's motors were run from a pair of switching torque control amplifiers that were interfaced, along with the encoders, to a multifunctional I/O card on a PC using some simple logic to implement safety features. Software control of the robot was achieved by means of a control loop running at 1,000 Hz, in which position and force were measured and desired output force was set. For further technical details, see (Howard et al. 2009). Participants held the handle of the vBOT that constrained hand movements to the horizontal plane. Using a projection system, we overlaid virtual visual feedback into the plane of the movement. Each of the vBOT handles controlled the

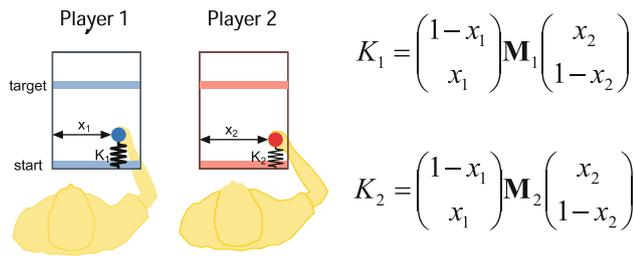


Fig. 1 Experimental setup. Player 1 and Player 2 each control a cursor from a start bar to a target bar. The movement can be chosen anywhere along the x -axes. However, the forces that resist the players' forward movement are given by spring constants whose stiffness depend both on the x -position of Player 1 and on the x -position of Player 2

position of a circular cursor (gray color, radius 1.25 cm) in one half of the workspace (Fig. 1). The position of the cursor was updated continuously throughout the trial.

Task design

At the beginning of each trial, 15-cm-wide starting bars were displayed on each side of the workspace (see Fig. 1). Participants were first required to place their cursor stationary within their respective starting bars, at which point two 15-cm target bars appeared at a distance randomly drawn from the uniform distribution between 5 and 20 cm (the same distance for both players on each trial). Each participant's task was to move their cursor to their respective target bar. This required them to make a forward movement (y -direction) to touch the target bar within a time window of 1,500 ms. The participants were free to touch it anywhere along its width (the robot simulated walls which prevented participants moving further laterally

than the width of the bar). During the movement, both players experienced a simulated one-dimensional spring, with origin on the starting bar, which produced a force on their hand in the negative y -direction thereby resisting their forward motion. The participants experienced forces $F_1 = -K_1 y_1$ and $F_2 = -K_2 y_2$ resisting their forward motion in which y_1 and y_2 are the y -distances of player 1 and 2's hands from the starting bar, respectively. The spring constants K_1 and K_2 depended on the lateral positions x_1 and x_2 of both players, where x corresponds to a normalized lateral deviation ranging between 0 and 1, such that the edges of the target bar (-7.5 and $+7.5$ cm) correspond to values 0 and 1. On each trial, the assignment of whether the left hand side or the right hand side of each target corresponded to 0 or 1 was randomized. Payoffs for intermediate lateral positions were obtained by bilinear interpolation:

$$K_i = \alpha \begin{pmatrix} 1-x_1 \\ x_1 \end{pmatrix}^T \begin{pmatrix} M_{0,0}^{(i)} & M_{1,0}^{(i)} \\ M_{0,1}^{(i)} & M_{1,1}^{(i)} \end{pmatrix} \begin{pmatrix} x_2 \\ 1-x_2 \end{pmatrix}$$

The 2-by-2 payoff matrices $M^{(1)}$ and $M^{(2)}$ defined the boundary values for the spring constants for players 1 and 2 at the extremes of the $x_1 x_2$ -space. That is $M_{j,k}^{(i)}$ is the payoff received by player i if player 1 takes action $x_1 = j$ and player 2 takes action $x_1 = k$. The scaling parameter α was constant throughout the experiment at 0.19 N/cm. Each game is defined by a different set of payoff matrices given by

Coordination game	$M^{(1)} = \begin{pmatrix} 0.1 & 1 \\ 1 & 0.1 \end{pmatrix}$	and	$M^{(2)} = \begin{pmatrix} 0.1 & 1 \\ 1 & 0.1 \end{pmatrix}$,
Stag hunt	$M^{(1)} = \begin{pmatrix} 0.1 & 1 \\ 0.5 & 0.5 \end{pmatrix}$	and	$M^{(2)} = \begin{pmatrix} 0.1 & 0.5 \\ 1 & 0.5 \end{pmatrix}$,
Chicken	$M^{(1)} = \begin{pmatrix} 0.5 & 0.7 \\ 0.1 & 1 \end{pmatrix}$	and	$M^{(2)} = \begin{pmatrix} 0.5 & 0.1 \\ 0.7 & 1 \end{pmatrix}$,
Battle of sexes	$M^{(1)} = \begin{pmatrix} 0.1 & 1 \\ 1 & 0.5 \end{pmatrix}$	and	$M^{(2)} = \begin{pmatrix} 0.5 & 1 \\ 1 & 0.1 \end{pmatrix}$.

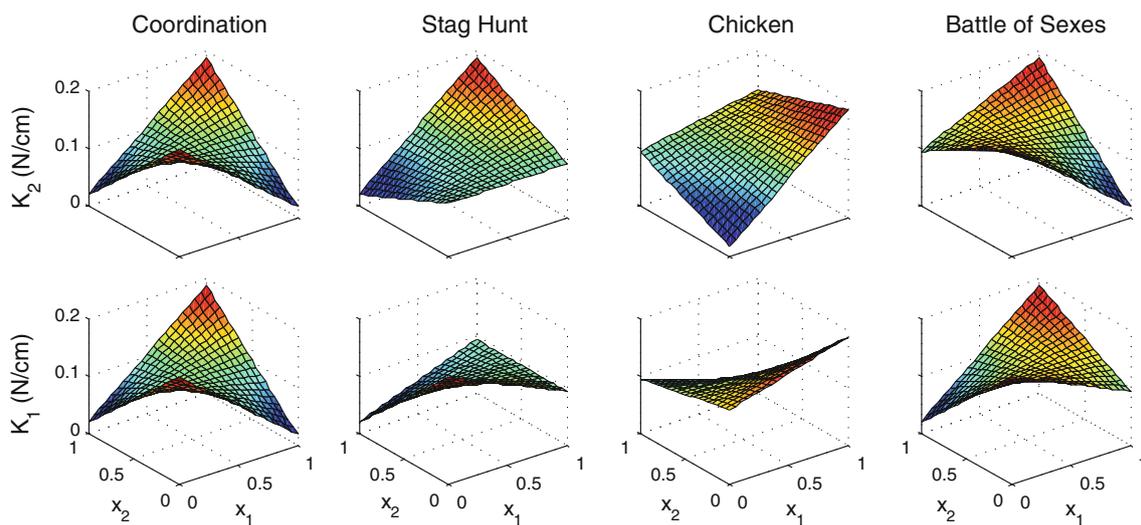


Fig. 2 Dependence of spring constants on the players' cursor positions. The four different games were associated with four different pairs of payoff matrices that induced different landscapes

of spring constants to generate spring-like forces that resisted players' forward movements. Regions with low spring constants are more desirable for players, as they are associated with less movement effort

The payoff landscapes resulting from these payoff matrices are plotted in Fig. 2. The pure Nash equilibria in these games are [right, right] or [left, left] in the choosing sides game, [stag, stag] and [hare, hare] in the stag hunt game, [swerve, stay] and [stay, swerve] in the chicken game, and [opera, opera] and [football, football] in the battle of sexes. Each pair of subjects performed all four types of games. The type of game changed every 20 trials. However, since within a set of 20 trials allocation of x_i being 0 and 1 was randomized to the left and right edges of the target, subjects had to coordinate their actions anew in each trial. This trial-by-trial randomization of the action allocation is equivalent to permuting the payoff matrix on each trial. All participants performed at least 10 blocks, but maximally 15 blocks of each game (i.e., at least $10 \times 20 = 200$ trials, maximum 300 trials per game type). Subjects were instructed to move the cursor to the target bar in such a way so as to minimize the resistive forces that they experienced. Furthermore, subjects were not permitted to communicate with each other during the experiment.

Data analysis

The final x -position of each player was categorized as “Action 0” or “Action 1” depending on whether $x < 0.5$ or $x > 0.5$. Coordination was considered successful when the “Actions” were chosen according to the action choices corresponding to the pure Nash equilibria, that is, when the endpoints x_1 and x_2 were in one of the two quadrants of the x_1x_2 -plane containing the pure Nash equilibrium payoffs of the matrix game. If the endpoints were in one of the other two quadrants, then the trial was considered to be miscoordinated. For the analysis of mutual information and joint entropy, we discretized the x -space in bins so as to form three sets S_1 , S_2 , and S_3 corresponding to the intervals $0 \leq x \leq 0.25$, $0.25 < x < 0.75$, and $0.75 \leq x \leq 1$, respectively. The discretization was chosen this way to distinguish the “unspecific” middle set of the x -space ($0.25 < x < 0.75$) from the borders of the x -space corresponding to “Action 0” and “Action 1,” respectively. Trajectories were evaluated at 10 points equidistant in time through the movement and both the mutual information and the joint entropy were computed over ensembles of these trajectory points that belonged to trajectories with either successful or unsuccessful coordination at the end of the trial. The mutual information

$$MI(x_1, x_2) = \sum_{i,j=1}^3 p(x_1 \in S_i, x_2 \in S_j) \times \log_2 \left(\frac{p(x_1 \in S_i, x_2 \in S_j)}{p(x_1 \in S_i)p(x_2 \in S_j)} \right)$$

measures the dependence between the two variables x_1 and x_2 (with a lower bound of 0 implying independence). $MI(x_1, x_2)$ measures how well the position x_1 can be predicted from knowing the position x_2 in the ensemble of successful (or unsuccessful) coordination trials. It therefore expresses the stereotypy of the x_1x_2 relationship over different trials. Importantly, as the x_1x_2 -quadrants for coordinated and miscoordinated trials have the same area there is no a priori reason for there to be a difference in mutual information on coordinated or miscoordinated trials. That is perfect miscoordination and perfect coordination (with all reaches ending in S_1 or S_3) would have identical high mutual information at the end of the movement. Therefore, mutual information provides information as to how related the two participants movements are at different times into a trial for different final coordination types. The joint entropy

$$H(x_1, x_2) = - \sum_{i,j=1}^3 p(x_1 \in S_i, x_2 \in S_j) \times \log_2 p(x_1 \in S_i, x_2 \in S_j)$$

measures how spread out the distributions are over x_1x_2 across trials. As we discretized, the space into three bins there is a maximum joint entropy of 3.2 bits. The distributions over x_1x_2 at the different time points from which the mutual information and joint entropy have been computed can be seen in Supplementary Figures S1–S3. The empirical estimates of mutual information and joint entropy were computed using a package described in (Peng et al. 2005).

Computational model

We developed a computational model for two-player co-adaptation with the aim of qualitatively replicating some of the characteristic features of successful coordination. The model is a highly simplified abstraction of the task which is not intended to fit all features of the data but to reproduce some key qualitative features of the dataset. Conceptually, it considers co-adaptation as a diffusion process through a “choice space” and abstracts away from the fact that movements had to be made in order to indicate choices or that there might be biomechanical constraints. In the model, co-adaptation is simulated by two point masses undergoing a one-dimensional diffusion through the payoff landscape. Each particle might be thought of as representing a player. Diffusion processes can be described both at the level of the deterministically evolving time-dependent probability density (Fokker–Planck-equation) and at the level of the stochastically moving single particle (Langevin equation).

The Langevin equation for each particle i can be written as $m \frac{dv_i}{dt} = G_i(x_1, x_2) - \gamma v_i + \sigma \eta_i$, where $m = 0.1$ is the point mass, v_i its instantaneous velocity, $\gamma = 0.001$ is a viscous constant capturing dissipation, $G_i(x_1, x_2)$ is a notional force that corresponds to the gradient of the payoff landscape with respect to changes in that player's action: $G_i(x_1, x_2) = -\frac{\partial}{\partial x_i} K_i, \eta_i$ is a normally distributed random variable and $\sigma^2 = 15$. Note that the mass m and the speed v do not refer to the actual movement process, but capture the inertia and the speed of the adaptation process. For each game, we simulated 10,000 trajectories with 50 time steps each (with discretized time step 0.01). The x -values were constrained to lie within the unit interval by reflecting the velocities when hitting the boundaries of 0 and 1. To prevent cyclical bouncing of the walls, reflections were simulated as inelastic reflections, where the outgoing speed was only 30% of the ingoing speed (inelastic boundary conditions). The initial x -values were sampled from a distribution that was piecewise uniform (x -values between $0.25 < x < 0.75$ had double the probability than x -values with either $0 \leq x \leq 0.25$ or $0.75 \leq x \leq 1$). This initialization roughly corresponded to the average initial distribution found in our experiments. The initial velocities were set to 0. All the results were computed from the simulated x -values in the same way as the experimental data.

Results

In order to assess the endpoint distribution of the games, we defined the endpoint of a trial as the time when both cursors had crossed the target bar for the first time. Thus, an endpoint of a game can be plotted as a point in the plane spanned by the final decision of each player. Figure 3a shows the endpoint distributions of the action choices of all pairs of players in all trials of the four games. The endpoint distributions in all games deviate significantly from the initial distribution (two-sample Kolmogorov–Smirnov test, $P < 0.01$), which shows that subjects evolved clear preferences for certain action combinations during the trial that differed from the starting positions. We pooled all endpoints from all blocks and trials of the games, since performance was roughly stationary both over trials within a block of twenty trials as well as over blocks of trials (Supplementary Figure S4). This absence of over-trial learning was a consequence of the trial-by-trial randomization of the payoff matrices and allowed us to focus on within-trial co-adaptation during individual movements. In order to compare our experimental results to the theoretical predictions, we categorized subjects' continuous choices into four quadrants corresponding to the discrete action pairs in the classic discrete game (Fig. 3b). In the three

games “choosing sides,” “stag hunt,” and “battle of sexes” the two pure Nash equilibria had the highest probability (Fig. 3c). In these games, the coordinated solution was significantly more frequent than the miscoordinated solution ($P < 0.01$ in all three games, Wilcoxon ranksum test over the six different subject pairs). Thus, the probability of having a coordinated solution increased from chance level at the beginning of the movements to a significantly elevated number of coordinated trials at the end of the movements (Fig. 3d). This result implies that coordination was achieved within individual trials through a co-adaptation process between the two actors.

In the chicken game, the two pure Nash equilibria [swerve, stay] and [stay, swerve] have a clearly increased probability compared to the worst case outcome of [stay, stay]. However, the “miscoordinated” solution [swerve, swerve] appears to have a comparable probability to the coordinated solutions of the pure equilibria. This might reflect an interesting property of this game, because, unlike the other games, the two pure Nash equilibria do not constitute an evolutionarily stable strategy (Smith and Price 1973; Smith 1982; Houston and McNamara 1991). Evolutionarily stable strategies are a refinement of the Nash equilibrium concept, i.e., every evolutionarily stable strategy constitutes a Nash solution, but not every Nash solution is “evolutionary” stable. In particular, Nash solutions that are evolutionarily stable are stable to perturbations, because any alteration of the equilibrium strategies around the equilibrium point leads to a strictly lower payoff. Importantly, in the chicken game, the evolutionarily stable strategy is given by the mixed Nash equilibrium, where each player stays or swerves with a certain probability. Depending on the perceived utilities of the experienced forces, the evolutionarily stable strategy might explain the observed response patterns. Thus, the results of this chicken game might also be compatible with single-trial coordination.

In order to elucidate the features that characterize successful single-trial coordination, we investigated kinematic differences between trials with successful coordination and trials with coordination failure (miscoordination). First, we assessed the statistical dependence between the two movement trajectories by computing the mutual information and the joint entropy over a normalized within-trial time course of the trajectories. We found higher joint entropy during the movement for trials with successful coordination compared to trials with coordination failure for the coordination game, stag hunt, and battle of sexes ($P < 0.05$ for each of the three games, Wilcoxon ranksum test on average joint entropy of sixth to eighth time point)—see Fig. 4a (upper panels). This suggests that the ensemble of successful coordination trajectories has more random variation. We also computed the mutual information between the two players over the same normalized

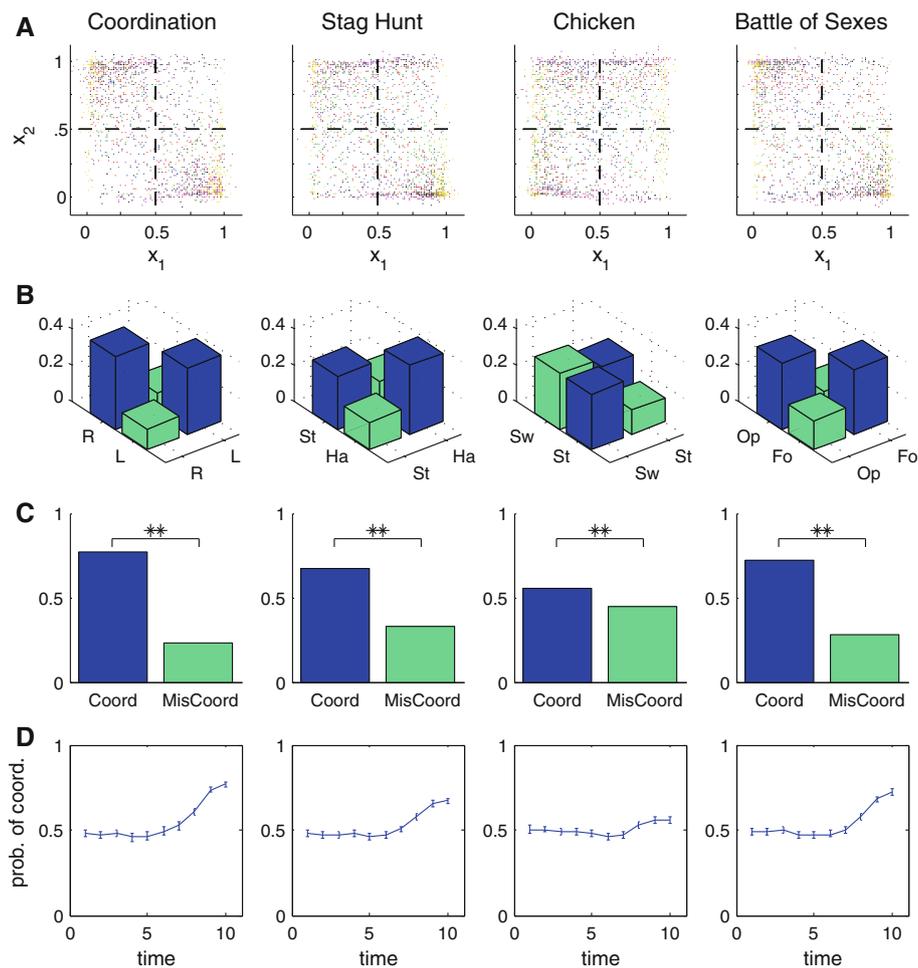


Fig. 3 **a** Endpoint distributions of the four different games for all players. The x_1 -position corresponds to the final position chosen by Player 1 and the x_2 -position to the final position of Player 2. The six different colors correspond to the six different subject pairs. **b** The distribution of the endpoints was binned into the four quadrants corresponding to the two times two actions of the respective discrete game. The quadrants in the coordination game correspond to the actions left (*L*) and right (*R*). The quadrants in the stag hunt game correspond to the actions stag (*St*) and hare (*Ha*). The quadrants in the chicken game correspond to the actions swerve (*Sw*) and stay (*St*).

The quadrants in the battle of sexes game correspond to football (*Fo*) and opera (*Op*). **c** Histogram over coordinated versus miscoordinated solutions. In all games and conditions, there was a significantly higher frequency of the coordinated outcomes. The *two stars* indicate a significance level of $P < 0.01$ in a Wilcoxon ranksum test over the six different games. **d** Temporal evolution of coordination probability. Trajectories were binned into 10 equidistant points. The coordination probability is given by the fraction of trajectory points that lie in the two quadrants corresponding to successful coordination. The *error bars* are obtained by bootstrapping

time windows (Fig. 4b). At the beginning of the trial, the mutual information was close to zero in all games and increased only toward the end of the trial. However, in coordinated trials the final mutual information was elevated compared to the final mutual information in miscoordinated trials ($P < 0.05$ in each of the four games, Wilcoxon ranksum test). This means that the final positions in coordinated trials tend to cluster together, whereas the final positions in miscoordinated trials are more scattered. Accordingly, the two actors share more information at the end of trials with successful coordination.

Since the pure Nash equilibria can be considered as attractors in the payoff landscape, it is reasonable to expect

that the initial position has an influence on the probability of reaching a particular Nash equilibrium, depending on whether the initial point is in the basin of attraction or not. We therefore examined the effect of the initial configuration on the probability of coordination success at the end of the movements. Figure 4c shows a histogram of the probability of successful coordination depending on the initial positions of the players. In particular, it can be seen that initial positions that are close to the pure Nash equilibria lead to an increased probability of coordination success at the end of the movements (Fig. 4d). Independent of the starting position, coordination success at the end of the movement is always greater than chance level. However,

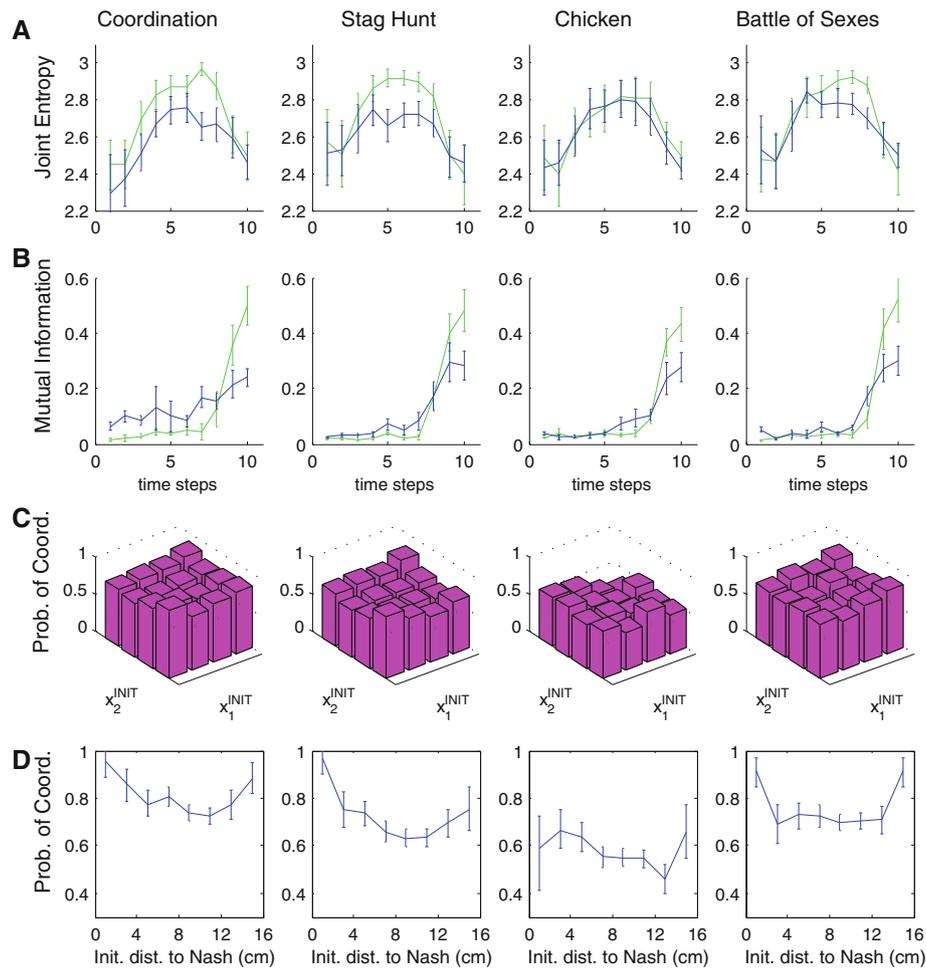


Fig. 4 **a, b** Joint entropy and mutual information between the distribution of positions of Player 1 and Player 2. All trajectories were discretized into 10 equidistant points and the positions of these points were categorized (see “Methods” for details). The average over coordinated trajectories is shown in blue and the average over miscoordinated trajectories is shown in green. During the movement, the joint entropy is increased for coordinated solutions in the

coordination, stag hunt, and battle of sexes game. The mutual information between the two players is always elevated at the end of the movement for all games compared to miscoordinated trials. **c, d** Dependence of coordination probability on initial positions of Player 1 and Player 2. For initial positions close to the corners of the workspace, the coordination probability at the end of the trial is increased. The *error bars* in **(d)** are obtained through bootstrapping

successful coordination is more likely to occur when players accidentally start out in favorable positions.

To further assess the effect of within-trial trajectory variation on coordination success, we examined the timing at which subjects’ trajectories showed little change from the final decision given by the endpoint when both players crossed the target bar. Specifically, we determined the earliest time at which each player’s trajectory stayed within 2 cm of their final lateral endpoint. This allowed us to assess whether players converged to their final decision early on (early convergence time) or whether the final position was chosen more abruptly (late convergence time). The time difference in convergence between the two players could then be determined as the absolute difference between their convergence times. We found that in successful coordination trials the convergence time difference

between players was significantly elevated compared to miscoordinated trials ($P < 0.02$, Wilcoxon ranksum test for equal medians of convergence time difference for all game types). In particular, we found that this difference arises from an early convergence of the first player to the final solution ($P < 0.02$, Wilcoxon ranksum test for equal medians of early convergence time for all game types) and is not due to a later convergence of the other player or an increase in total trial time (Fig. 5). This suggests that successful coordination is in general more likely to occur when one of the players converges early to their position.

Finally, we devised a simple computational model to investigate qualitatively a potential computational basis of some of the features characteristic for within-trial coordination (see “Methods” for details). The computational model can be thought of as a diffusion process through the

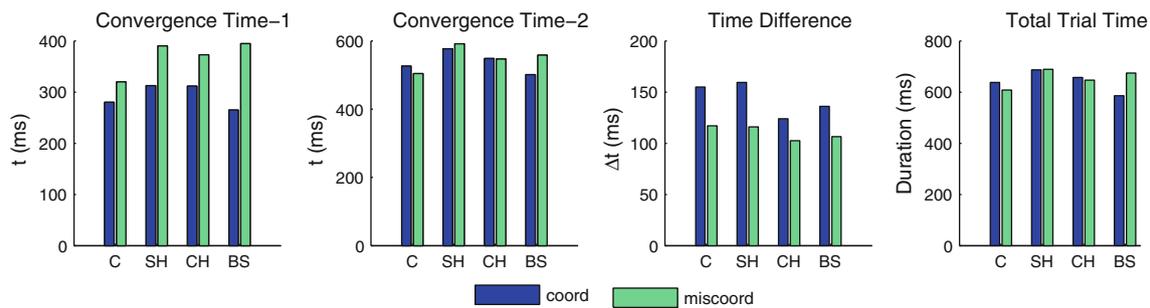


Fig. 5 Difference in convergence times. We computed the time of convergence to a particular choice as the time after which a player did not leave a cone of 2 cm width. In coordinated trials (*blue*), the difference of the convergence times between the two players was always significantly elevated compared to miscoordinated trials (*green*), and the time the first player converged to the final solution (Convergence Time-1) was significantly lower compared to

miscoordinated trials. This was true for the coordination game (C), the stag hunt (SH), the chicken game (CH), and the battle of sexes game (BS). In contrast, the later convergence time and the total trial time was not significantly elevated ($P > 0.1$, Wilcoxon ranksum test for equal medians), except for the Battle of Sexes game. Yet, even in the Battle of Sexes game, the effect of the early convergence time is much stronger than the effect of the later convergence time

payoff landscape and is a standard model in physics. In a physical system, the payoff would correspond to an energy surface and the players could be thought of as particles moving stochastically in this energy landscape, trying to move downhill (lower forces and hence higher payoff). Diffusion models have also been previously applied to understand coordination in iterative classic coordination games (Crawford 1995). In our simulations, the initial positions of the players were drawn from a distribution similar to the initial distribution observed experimentally. We then simulated diffusion in the payoff landscape, assuming that each player estimates a noisy version of how the spring stiffness varies with their lateral movement (local payoff gradient) and tries to move downhill. Importantly, there are no lateral forces in the experiment, which means the modeled “downhill movement” does not correspond to the robot forces that are pushing against the subjects’ forward motion, but represent the subject’s voluntary choice following a noisy gradient in spring-constant space (compare Fig. 2). This model allows us then to determine the final distribution of positions of each player (Fig. 6, top row). The model captures the increased relative occurrence of the Nash equilibria by showing that coordinated solutions occur much more frequently than miscoordinated solutions (Fig. 6, second row). The model accounts for the most frequent pattern of coordination in all games, with the exception of the increased probability of the experimentally observed [swerve, swerve] solution in the chicken game. As observed in the experimental data, the model also accounts for an increased joint entropy during the movement between the positions of the players in successful coordination trials as opposed to miscoordinated trials (Fig. 6, third row). The model also captures that the mutual information between the two players’ position should be increased at the end of the trial in case of successful coordination (Fig. 6, fourth row). Finally, the

model captures the increased difference in decision time in successful coordination trials compared to miscoordinated trials (Fig. 6, bottom row). However, there are also important features that are not captured by the diffusion model that are discussed below.

Discussion

In this study, we examined how coordination arises between two coupled motor systems that interact in continuous single-trial games with multiple Nash equilibria. Crucially, the presence of multiple equilibria creates the problem of equilibrium selection where players have to coordinate their actions in order to achieve low-cost outcomes. In contrast to discrete, one-shot coordination games, we found that subjects were able to coordinate within individual trials in our motor games. As the payoff matrix was permuted on each trial, finding a coordinated solution is only possible through a continuous co-adaptation process within a trial, in which the two players may share information. Such continuous within-trial adaptation processes are typical of motor interactions (Braun et al. 2009a). We tested subjects on four different coordination games to find common principles of successful coordination that govern all four different types of two-player motor interactions. Our aim was to use these four different games to confirm the generality of our findings rather than to investigate differences between the games. Successful coordination trials were characterized by several features: increased mutual information between the final part of the movements of the two players, an increase in joint entropy during the movements, and an early convergence to the final position by one of the players. Moreover, we found that the probability of coordination depended on the initial position of the two players, in particular on their distance to

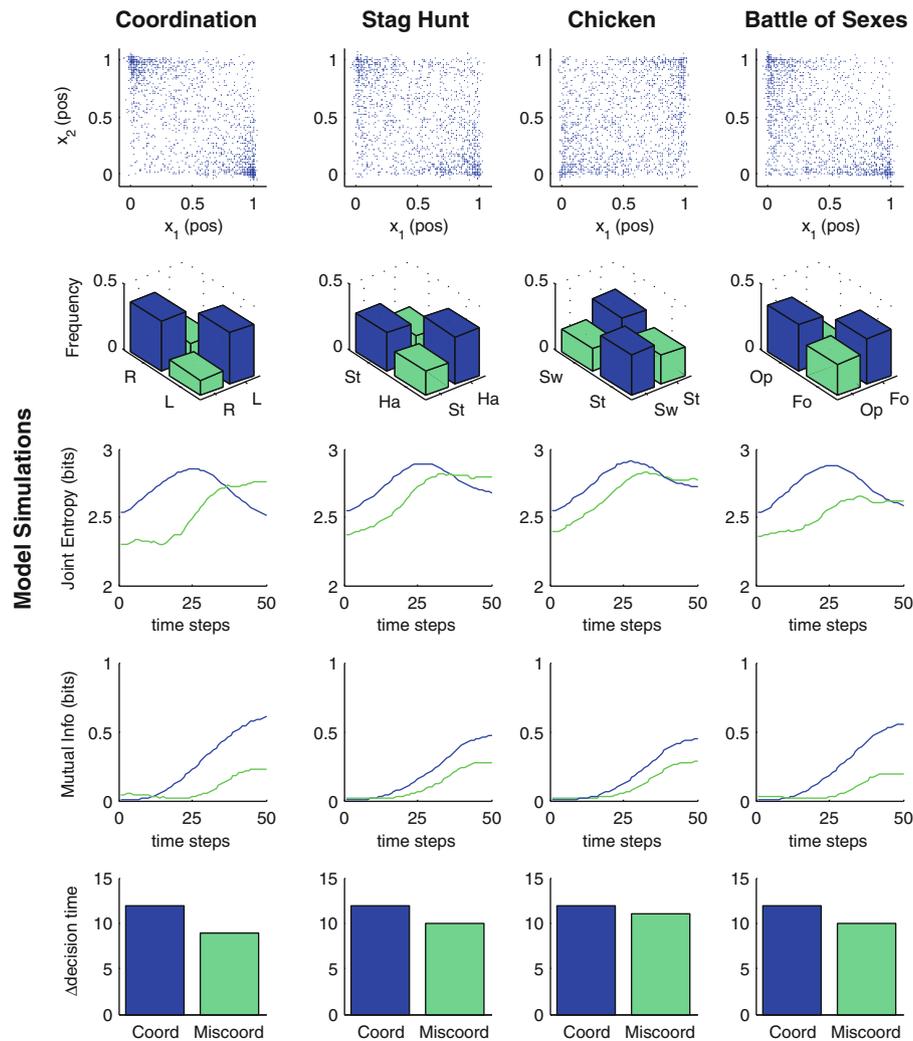


Fig. 6 Model simulations. (Row 1) Endpoint distribution of the diffusion process for the different games after 50 time steps. Thousand eight hundred draws are shown. (Row 2) Histogram over simulated endpoints. The histograms are similar to the experimental ones—compare Fig. 3. Analogous to the experimental data, endpoints can be classified as coordinated and miscoordinated. (Row 3) Joint entropy of the position distributions of player 1 and player 2. During the movement, the joint entropy is elevated in successful coordination trials (blue) compared to miscoordinated trials (green). This is a similar pattern observed in the experimental data—compare Fig. 4.

(Row 4) Mutual information between the position distributions of player 1 and player 2. At the end of the movement, the mutual information is elevated in successful coordination trials (blue) compared to miscoordinated trials (green). The same pattern is observed in the experimental data—compare Fig. 4. (Row 5) Temporal difference in convergence time. In coordinated trials, the difference in time when the players converge to their final position is increased compared to miscoordinated trials. The same pattern is observed in the experimental data—compare Fig. 5

one of the pure Nash equilibria. Our results suggest that coordination is achieved by a favorable starting position, a stereotypical coordination pattern, and an early convergence to a final position by one of the players.

We used a model to explore whether features seen in our results could arise from a simple diffusion model in which each subject samples the noisy payoff gradient with respect to their movement and tries to move downhill in the payoff landscape thereby reducing effort. We show that the features of joint entropy, mutual information, and convergence time difference can be reproduced qualitatively by

this model. However, there were also important discrepancies. First, the shape of the joint entropy functions does not correspond to the experimentally observed inverse-U shape for miscoordinated trials. Second, the mutual information increases steadily over the course of the trial and not abruptly at the end of the movement, as is observed experimentally (compare Figs. 3 and 6). Third, the model predicts that the probability of coordination decreases monotonically with the initial distance from the Nash equilibria (compare Supplementary Figure S5), whereas in the experimental data, the probability increases again at the

border of maximum distance (compare Fig. 4). First, it needs to be noted that the model completely ignores any movement dynamics and is only concerned with the abstract process of within-trial co-adaptation. Furthermore, in the experiment subjects have to sample the force landscape and do not know the infinitesimal gradient. Despite these shortcomings this simple model shows that coordination can in principle arise through gradient-descent-like co-adaptation and it reproduces a number of characteristic properties of two-player coordination. In the future, more complex models that take subjects' movement properties into account have to be developed to better understand the co-adaptation processes during two-player motor interactions.

In a previous study (Braun et al. 2009b), we have shown that two-player interactions in a continuous motor game version of the Prisoner's Dilemma can be understood by the normative concept of Nash equilibria. However, in the Prisoner's Dilemma the Nash solution is unique and corresponds to a minimax-solution, which means that each player could select the Nash solution independently and thus minimize the maximum effort. In the current study, each game had several equilibria giving rise to the novel problem of equilibrium selection. In these coordination games, it is no longer possible for each player to select their optimal strategy independently. Instead, players have to coordinate their actions in order to achieve optimal Nash solutions. They have to explore actively the payoff landscape in order to converge jointly to a coordinated solution. Our results suggest that this coordination arises naturally as the behavior of the actors gets increasingly correlated through a process of co-adaptation.

In the game-theoretic literature, diffusion processes have been suggested previously to model learning and co-adaptation in games (Foster and Young 1990; Fudenberg and Harris 1992; Kandori et al. 1993; Binmore et al. 1995; Binmore and Samuelson 1999), including the special case of coordination games (Crawford 1995). In particular, in evolutionary game theory, adaptation processes are often modeled as diffusion-like processes based on so-called replicator equations (Weibull 1995; Ruijgrok and Ruijgrok 2005; Chalub and Souza 2009). These replicator equations can be thought of as population equations that facilitate "survival" of useful strategies and leave unsuccessful strategies to perish. Solutions of the replicator equations correspond to Nash solutions that are stable to perturbations, that is, they are evolutionarily stable strategies (Smith 1982). In three of our motor game experiments the predicted pure Nash solutions corresponded to evolutionarily stable strategies. However, in the chicken game, this was not the case, and the evolutionarily stable strategy is a mixed Nash equilibrium. Interestingly, this is also the game where our model predictions were the poorest. In the

future, it might therefore be interesting to investigate whether the concept of evolutionarily stable strategies might be even more relevant to understand multi-player motor interactions than the concept of Nash equilibria.

Diffusion processes have also been used extensively to model decision-making in psychology and neurobiology (Busemeyer and Townsend 1993; Busemeyer and Diedrich 2002; Smith and Ratcliff 2004; Gold and Shadlen 2007). In such models, evidence is accumulated before a decision boundary is reached and the appropriate action executed. In our study, the evidence accumulation process happens during the movement in a continuous fashion before the final decision at the target bar is taken. Therefore, the co-adaptation process for coordination might also be conceived as a continuous decision-making process linking the neurobiology of decision-making to the neurobiology of social interactions. Previous studies on the neuroscience of social interactions have suggested common principles for motor control and social interactions, in particular the usage of forward internal models for predictive control and action understanding (Wolpert et al. 2003). In our task, we have not found evidence for model-based predictive control, and our model simulations show that single-trial coordination can be achieved by co-adaptive processes performing essentially gradient descent on a force landscape. This might be a consequence of our task design, as the force landscape changed every trial such that subjects could not predict the payoffs. In this task, participants were also not aware of the consequence of their actions on the other player, since each player could not feel the force feedback given to the other player. However, when we conducted the same experiment and gave visual feedback of the forces experienced by the other player, the results remained essentially unchanged (see Supplementary Figure S6). It could be the case, however, that subjects simply ignored this feedback signal, even though it was explained to them before the start of the experiments. Our results suggest therefore that single-trial coordination in continuous motor tasks can be achieved in the absence of predictive control. However, tasks in which the payoff landscape is constant over trials, and therefore learnable, might facilitate the formation of predictive internal models for two-person motor interactions.

In previous studies, it has been shown that natural pattern of coordination can arise between participants when provided with feedback of the other participant, such as the synchronization of gait patterns—for a review see (Kelso 1995). This coordination is typically explained as the consequence of stable configurations of nonlinearly interacting components (Haken et al. 1985). Game-theoretic models of adaptation are very similar to these dynamical systems models in the sense that they also converge to stable solutions given by fix points or other attractors. The

force landscape of our coordination landscape can be conceived as an attractor landscape where the pure Nash equilibria correspond to point attractors. Consequently, the convergence to a particular equilibrium depends on the initial position starting out in a particular basin of attraction of one of the equilibria. The interesting difference to purely dynamical systems models is that game-theoretic models allow making a connection to normative optimality models of acting. In our case, that means that the Nash equilibria are not only attractors in a dynamical system, but they also represent optimal solutions to a game-theoretic optimization problem. Consequently, movements are not only interpretable in terms of an abstract attractor landscape but also in terms of behaviorally relevant cost parameters such as movement accuracy or energy consumption. Thus, the two perspectives of dynamical systems and optimality models of acting are not only compatible but also complement each other (Schaal et al. 2007). In summary, our study has three main novel findings. First, successful coordination is facilitated by a favorable starting state of both players. Second, coordination reflects a more stereotypical coordination pattern between players when compared to miscoordinated trials. Finally, when one player reaches an early convergence time relative to the other player, this increases the chance of successful coordination. We also show that these features can be qualitatively fit by a simple diffusion model of coordination.

Acknowledgments We thank Lars Pastewka for discussions and Ian Howard and James Ingram for technical assistance. Financial Support was provided by the Wellcome Trust and European project SENS-OPAC IST-2005 028056.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Aumann R (1974) Subjectivity and correlation in randomized strategies. *J Math Econ* 1:67–96
- Aumann R (1987) Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55:1–18
- Binmore K, Samuelson L (1999) Evolutionary drift and equilibrium selection. *Rev Econ Stud* 66:363–393
- Binmore K, Samuelson L, Vaughan R (1995) Musical chairs: modeling noisy evolution. *Games Econ Behav* 11:1–35
- Braun DA, Aertsen A, Wolpert DM, Mehring C (2009a) Learning optimal adaptation strategies in unpredictable motor tasks. *J Neurosci* 29:6472–6478
- Braun DA, Ortega PA, Wolpert DM (2009b) Nash equilibria in multi-agent motor interactions. *PLoS Comput Biol* 5:e1000468
- Busemeyer JR, Diedrich A (2002) Survey of decision field theory. *Mathematical Social Sciences* 43:345–370
- Busemeyer JR, Townsend JT (1993) Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol Rev* 100:432–459
- Chalkiadakis G, Boutilier C (2003) Coordination and multiagent reinforcement learning: a Bayesian approach. In: Proceedings of the second international conference on autonomous agents and multiagent systems, Melbourne
- Chalub FA, Souza MO (2009) From discrete to continuous evolution models: a unifying approach to drift-diffusion and replicator dynamics. *Theor Popul Biol* 76:268–277
- Crawford VP (1995) Adaptive dynamics in coordination games. *Econometrica* 63:103–143
- Foster D, Young P (1990) Stochastic evolutionary game dynamics. *Theor Popul Biol* 38:219–232
- Fudenberg D, Harris C (1992) Evolutionary dynamics with aggregate shocks. *J Econ Theory* 57:420–441
- Fudenberg D, Tirole J (1991) *Game theory*. MIT Press, Cambridge
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574
- Greenwald A, Hall K (2003) Correlated Q learning. In: Proceedings of the twentieth international conference on machine learning, Washington
- Haken H, Kelso JA, Bunz H (1985) A theoretical model of phase transitions in human hand movements. *Biol Cybern* 51:347–356
- Houston AI, McNamara JM (1991) Evolutionary stable strategies in the repeated hawk-dove game. *Behav Ecol* 2:219–227
- Howard IS, Ingram JN, Wolpert DM (2009) A modular planar robotic manipulandum with end-point torque control. *J Neurosci Methods* 181:199–211
- Kandori M, Mailath G, Rob R (1993) Learning, mutation, and long run equilibria in games. *Econometrica* 61:29–56
- Kelso JAS (1995) *Dynamic patterns: the self-organization of brain and behavior*. MIT Press, Cambridge
- Nash JF (1950) Equilibrium points in N-Person games. *Proc Natl Acad Sci USA* 36:48–49
- Peng H, Long F, Ding C (2005) Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans Pattern Anal Mach Intell* 27:1226–1238
- Rousseau J-J, Bachofen B, Bernardi B (2008) *Discours sur l'origine et les fondements de l'inégalité parmi les hommes*. Flammarion, Paris
- Ruijgrok M, Ruijgrok T (2005) Replicator dynamics with mutations for games with a continuous strategy space. arXiv: nlin/0505032v2
- Schaal S, Mohajerian P, Ijspeert A (2007) Dynamics systems vs. optimal control—a unifying view. *Prog Brain Res* 165:425–445
- Smith JM (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge
- Smith MJ, Price GR (1973) The logic of animal conflict. *Nature* 246:15–18
- Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27:161–168
- Weibull JW (1995) *Evolutionary game theory*. MIT Press, Cambridge
- Wolpert DM, Doya K, Kawato M (2003) A unifying computational framework for motor control and social interaction. *Philos Trans R Soc Lond B Biol Sci* 358:593–602