



Research

Cite this article: van den Berg P, Weissing FJ. 2015 The importance of mechanisms for the evolution of cooperation. *Proc. R. Soc. B* **282**: 20151382.
<http://dx.doi.org/10.1098/rspb.2015.1382>

Received: 10 June 2015

Accepted: 9 July 2015

Subject Areas:

behaviour, evolution, theoretical biology

Keywords:

social behaviour, simulation model, Prisoner's Dilemma, Snowdrift game, neural network, non-equilibrium dynamics

Author for correspondence:

Franz J. Weissing

e-mail: f.j.weissing@rug.nl

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rspb.2015.1382> or via <http://rspb.royalsocietypublishing.org>.

The importance of mechanisms for the evolution of cooperation

Pieter van den Berg and Franz J. Weissing

Groningen Institute for Evolutionary Life Sciences, University of Groningen, Groningen 9747 AG, The Netherlands

Studies aimed at explaining the evolution of phenotypic traits have often solely focused on fitness considerations, ignoring underlying mechanisms. In recent years, there has been an increasing call for integrating mechanistic perspectives in evolutionary considerations, but it is not clear whether and how mechanisms affect the course and outcome of evolution. To study this, we compare four mechanistic implementations of two well-studied models for the evolution of cooperation, the Iterated Prisoner's Dilemma (IPD) game and the Iterated Snowdrift (ISD) game. Behavioural strategies are either implemented by a 1 : 1 genotype–phenotype mapping or by a simple neural network. Moreover, we consider two different scenarios for the effect of mutations. The same set of strategies is feasible in all four implementations, but the probability that a given strategy arises owing to mutation is largely dependent on the behavioural and genetic architecture. Our individual-based simulations show that this has major implications for the evolutionary outcome. In the ISD, different evolutionarily stable strategies are predominant in the four implementations, while in the IPD each implementation creates a characteristic dynamical pattern. As a consequence, the evolved average level of cooperation is also strongly dependent on the underlying mechanism. We argue that our findings are of general relevance for the evolution of social behaviour, pleading for the integration of a mechanistic perspective in models of social evolution.

1. Introduction

There is a long tradition in biology of separating proximate and ultimate perspectives when explaining phenotypic variation [1,2]. The proximate perspective is concerned with the mechanisms that directly cause the phenotype (such as neurological and physiological processes), whereas the ultimate perspective is concerned with the emergence of the phenotype through (adaptive) evolution. In concordance with this traditional separation, knowledge about the specific mechanisms underlying phenotypes has long been regarded as inconsequential to the question of how phenotypes are shaped by evolution. Accordingly, evolutionary biologists have a strong focus on the fitness consequences of phenotypic traits, thereby largely disregarding the underlying mechanisms. Conceptualization of evolution is often based on the implicit assumptions that genes interact in a simple way and that there is a one-to-one relationship between genotypes and phenotypes. These assumptions are convenient, as they allow a view of selection as a process directly acting on the genes in the 'gene pool' of a population. Although this view has already been criticized as 'beanbag genetics' more than 50 years ago [3], theoretical approaches to explaining the evolution of phenotypes with an explicit focus on mechanisms are not very prominent even today.

Verbal discussions of the importance of underlying mechanisms for the dynamics and outcomes of evolutionary processes started to emerge in the literature in the 1980s [4]. In particular, the influential book of John Maynard Smith and Eörs Szathmáry on the 'Major Transitions in Evolution' [5] clearly showed how crucial genetic and phenotypic architecture are for the course of evolution. This view is now firmly established in the field of 'evo-devo' [6,7], where the interplay between (developmental) mechanisms and evolution is at centre stage. Similarly, studies on gene-regulatory networks [8–10] have revealed that network topology strongly affects both the robustness and evolvability of living systems, while recent 'integrative' models [11–13] reveal that the mechanisms

underlying phenotypic responses can be important for a full understanding of eco-evolutionary processes.

In line with these general developments, there are now strong pleas [14–16] to apply ‘mechanistic thinking’ in evolutionary studies of animal and human behaviour as well. Yet, with some notable exceptions [17–22], models for the evolution of behaviour still tend to make the ‘least constraining’ assumptions on the genetic basis and the physiological and psychological processes underlying behaviour. When the direction and intensity of selection do not change in time and when there is a single optimal behaviour, this may not be problematic. In such a case, one would expect evolution to proceed towards the single optimum, regardless of underlying mechanisms. However, whenever there are multiple equilibria, the situation is no longer so straightforward. And even in relatively simple social contexts, the existence of multiple equilibria is the rule rather than the exception [23–25]. In other words, the question is not that much ‘which strategy is favoured by natural selection’ but rather ‘which equilibrium will be achieved in the course of evolution’ [26–28]. It is conceivable that, in such a context, the mechanisms underlying behaviour may be of evolutionary importance, because mechanisms can affect the probabilities with which phenotypes arise and, hence, the likelihood of alternative evolutionary trajectories.

Here, we study the evolution of behavioural strategies in two types of social interaction without clearly delineated optimal behaviour. Our question is whether, and to what extent, the mechanistic implementation of the available strategies affects the course and outcome of evolution. We consider two prototype models for the evolution of cooperation: the Iterated Prisoner’s Dilemma (IPD) game and the Iterated Snowdrift (ISD) game, which have been the subject of hundreds of earlier studies (IPD [24,29–34], ISD [33–37]). In both games, the players have to decide (repeatedly) on whether to cooperate or to defect. For both players, mutual cooperation is more profitable than mutual defection. However, mutual cooperation is not easy to achieve, as defection yields a higher pay-off than cooperation if the other player cooperates. The games differ in their assumption on whether defect (IPD) or cooperate (ISD) yields a higher pay-off against a defector. Following the traditions of evolutionary game theory ([28,38–45], but see [46]), studies of the evolution of strategies in these games have overwhelmingly assumed a one-to-one relationship between genotypes and strategies. Here, we contrast such a one-to-one implementation with a different implementation where selection does not directly act on strategies, but on the architecture (a simple neural network) underlying these strategies. In addition, we consider two genetic mechanisms that determine the probabilities with which the mutation of each strategy yields any other strategy. We will show that the evolutionary dynamics are strongly affected by both the genetic and the behavioural architecture and discuss how the different outcomes can be explained on the basis of the mutational distributions arising from the interplay between genetics and behavioural mechanisms.

2. The model

(a) Games and strategies

Throughout, we will consider a Prisoner’s Dilemma (PD) game and a Snowdrift game (SD) with the following pay-off matrices:

$$\text{PD:} \begin{pmatrix} 3 & 0 \\ 5 & 1 \end{pmatrix} \quad \text{and} \quad \text{SD:} \begin{pmatrix} 3 & 1 \\ 5 & 0 \end{pmatrix}.$$

The top and bottom rows give the pay-offs of cooperation and defection, respectively, both for when the opponent cooperates (first column) and defects (second column). In the PD, defection always yields a higher pay-off than cooperation, regardless of the action of the opponent. In the one-shot version of this game, mutual defection is therefore the only evolutionarily stable strategy (ESS). In the SD, the highest pay-off is always attained by choosing the opposite action than the opponent. In this case, the one-shot game has an ESS that is characterized by a mixture of cooperation and defection.

We consider iterated versions of both games, for which the determination of all ESSs is much less straightforward than for their one-shot counterparts (see the electronic supplementary material for a game-theoretical analysis). In our simulations, agents repeatedly interact for an indefinite period of rounds; after each round, the game is terminated with probability $1 - m$. The full strategy space of the iterated game is infinite-dimensional [30]. Here, we confine the strategy space by only allowing individuals to condition their behaviour on the outcome of the previous interaction round. As there are four possible interaction outcomes (mutual cooperation, mutual defection, and both combinations of cooperation and defection), and a strategy always prescribes one of two possible actions for each outcome (cooperation or defection), there are in total $2^4 = 16$ possible strategies (see table 1 for a complete list). We assume that individuals are not perfect; they make both perception errors (with probability ε_P , they misinterpret the behaviour of their opponent as the opposite behaviour) and implementation errors (with probability ε_I , they perform the opposite behaviour than dictated by their strategy).

(b) Behavioural and genetic architecture

Figure 1 shows a schematic of the behavioural and genetic architectures considered in this study. We consider a ‘1 : 1’ behavioural architecture and an artificial neural network (ANN) architecture, which can both realize the 16 possible strategies presented in table 1. The 1 : 1 architecture is the simplest possible architecture, in which behaviour for each of the four possible outcomes of the previous round is under the direct control of a single gene locus. Each of these four loci can only have two values: 1 (for cooperation) or 0 (for defection). In addition, a separate locus determines an individual’s behaviour in the first round; this locus can take on any value of the unit interval, which corresponds to the probability of cooperation in the first round.

In the neural network architecture ([47]; see figure 1 for a graphical representation; and see the electronic supplementary material for a more detailed explanation), behaviour is determined through a very simple underlying structure that translates an input (the behaviours of ‘self’ and ‘partner’ in the previous interaction round) into an output (cooperation or defection). There are two input nodes, one of which receives the previous own behaviour (0 for defection, 1 for cooperation), and the other receives the previous behaviour of the opponent. The input from both these nodes is fed into two ‘hidden layer’ nodes, multiplied by the weights of the connections between the nodes. Each hidden layer node has a threshold; if the summed input into a hidden layer node exceeds its threshold, its output equals 1, otherwise the output is 0. Both hidden layer nodes are connected to

Table 1. An overview of all possible pure strategies that condition their behaviour on the previous interaction. (The four columns on the left show whether the strategy cooperates (1) or defects (0), for each of the four possible outcomes of the previous interaction (from left to right: mutual cooperation, having cooperated while the opponent defected, having defected while the opponent cooperated, and mutual defection). The column on the right shows the name of the strategy that is used to refer to them in the main text. The two middle columns show the percentage of the genotype space that is associated with each strategy for the two different behavioural architectures. These percentages were obtained by generating a large number of genotypes (in the same way as generating a genotype through ‘entire-genome mutation’), and subsequently determining the strategy induced by each genotype (as explained in the electronic supplementary material).)

behaviour				percentage of genotype space		neural network	strategy name	strategy description
C	C	D	D	1 : 1 mapping				
C	D	C	D					
0	0	0	0	6.25	40.35	<i>ALLD</i>	always defects	
0	0	0	1	6.25	1.79	<i>desperate</i>	only cooperates after mutual defection	
0	0	1	0	6.25	1.75	<i>Acon-D</i>	anti-conventional, shifts after playing opposite of opponent, otherwise defects	
0	0	1	1	6.25	1.65	<i>inconsistent</i>	plays opposite of previous move	
0	1	0	0	6.25	1.75	<i>con-D</i>	conventional, stays after playing the opposite of opponent, otherwise defects	
0	1	0	1	6.25	1.65	<i>ATFT</i>	anti-tit for tat, plays opposite of opponent’s last move	
0	1	1	0	6.25	0.08	<i>APavlov</i>	win, shift; lose, stay	
0	1	1	1	6.25	0.98	<i>hopeless</i>	only defects after mutual cooperation	
1	0	0	0	6.25	0.98	<i>grim</i>	only cooperates after mutual cooperation	
1	0	0	1	6.25	0.08	<i>Pavlov</i>	win, stay; lose, shift	
1	0	1	0	6.25	1.65	<i>TFT</i>	tit for tat, copies opponent’s last move	
1	0	1	1	6.25	1.75	<i>MNG</i>	Mr Nice Guy, only defects after ‘being cheated’ (playing C while other plays D)	
1	1	0	0	6.25	1.65	<i>consistent</i>	repeats its own previous move	
1	1	0	1	6.25	1.75	<i>con-C</i>	conventional, stays after playing the opposite of opponent, otherwise cooperates	
1	1	1	0	6.25	1.79	<i>willing</i>	only defects after mutual defection	
1	1	1	1	6.25	40.35	<i>ALLC</i>	always cooperates	

the output node, which also has a threshold. If the total output from the hidden nodes (multiplied with the relevant connection weights) exceeds this threshold, the individual cooperates. If not, the individual defects. This way, six connection weights and three thresholds determine the strategy implemented by the network. Accordingly, the ANN is encoded by nine gene loci (that can take on any real value): one for each connection weight and one for each threshold. In addition, a tenth locus determines an individual’s behaviour in the first round (as in the 1 : 1 implementation).

Even under the highly simplifying assumptions on the strategy set and their underlying architectures, there are many ways to implement inheritance. For example, ploidy level and linkage patterns are of considerable importance for the genetic transmission of information on strategies in sexually reproducing organisms. To keep matters as simple as possible, we here only consider asexual populations of haploid individuals. In order to study the effect of genetic factors, we restrict attention to two different mutation regimes. In both regimes, the gene locus determining the behaviour in the first round mutates independently (with

probability μ_F), the mutational step size being drawn from a normal distribution with mean 0 and standard deviation σ_F . Under ‘per-locus mutation’, each of the other loci (four loci in case of the 1 : 1 architecture and nine in case of the ANN architecture) has a probability μ_L of giving rise to a mutation, independently of what is happening at the other loci. Under ‘entire-genome mutation’, a mutation event (occurring with probability μ_G) affects all these loci, that is, all these loci mutate at the same time. Under both mutation regimes, mutation is implemented as drawing a random number to replace the current value of the locus (in the 1 : 1 architecture, this is done by drawing 0 or 1 with equal probability; in the neural network architecture, by drawing a number from a normal distribution with mean 0 and standard deviation σ_N).

(c) Simulation set-up and parameters

We simulated a population of 1000 asexual haploid individuals, with discrete and non-overlapping generations. At the start of each generation, pairs of two individuals are

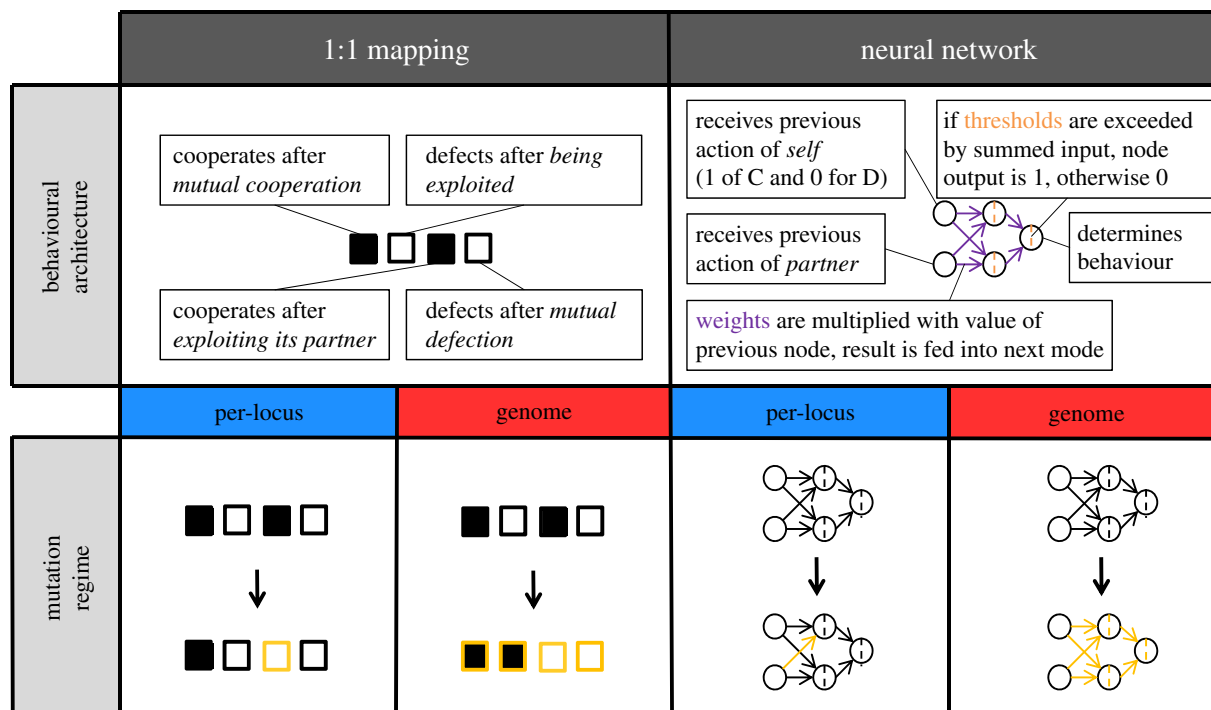


Figure 1. A schematic of the four implementations of the 16 strategies considered in this study. The top row shows illustrations of the two behavioural architectures. In the 1:1 architecture, individuals have four gene loci that each determine the behaviour (cooperate or defect) in a given round for one of the four possible outcomes of the previous round. These four loci are represented by boxes (in the example shown, black boxes represent cooperation and white boxes represent defection). In the neural network architecture, individuals have nine loci, determining the (continuous) values of six connection weights (purple) and three thresholds (orange). The network processes the input (the behaviour of 'self' or 'partner' in the previous round) into an output (cooperate or defect). In the bottom row, the two mutation regimes are illustrated for both behavioural architectures, representing the four implementations considered in this study. Under per-locus mutation, each locus mutates independently (illustrated by single loci turning yellow after the arrow). In case of whole-genome mutation, all loci mutate in the event of a mutation (illustrated by all loci turning yellow after the arrow).

formed at random. These pairs interact repeatedly, where a new round always starts with probability $m = 0.99$ (leading to an average interaction length of 100 rounds). In any given generation, all pairs play the same number of rounds. After the last round of each repeated interaction, individuals reproduce. The probability of reproducing is directly proportional to the pay-off individuals accumulate over the entire repeated interaction. Population size was kept constant. After reproduction, a new cycle starts.

At the beginning of each simulation, the loci of all individuals were initialized at random: initial values for the locus that determines the behaviour in the first round were drawn from a normal distribution with mean 0.5 and standard deviation 0.1; the four binary loci in the 1:1 architecture were assigned a 0 or a 1 with equal probability; and the nine loci encoding the connection weights and thresholds in the neural network architecture were assigned values that were drawn from a normal distribution with mean 0 and standard deviation σ_N . Each simulation was run for 100 000 generations. We ran 100 replicate simulations for all four combinations of the two behavioural architectures (1:1 and neural network) and the two mutation regimes (per-locus mutation and entire-genome mutation). Resulting cooperation levels and strategy frequencies were calculated by averaging over all interactions in the last generation of each simulation, and then averaging those averages over all replicates.

In all simulations reported here, the perception error ϵ_P and the implementation error ϵ_I were both set to 0.01; mutation probabilities (μ_L , μ_G and μ_F) were all set to

0.001, and mutational step sizes for all continuous loci (σ_N and σ_F) were set to 0.1. In the electronic supplementary material, we consider different values of these parameters in order to check for the robustness of our results.

3. Results

(a) Effect of architecture on the average cooperation level

We studied the evolution of cooperation in two games (the ISD and the IPD), with four different implementations (figure 1) of behavioural strategies, reflecting two scenarios concerning the underlying behavioural architecture (1:1 versus neural network), and two scenarios concerning the mutation regime (per-locus versus entire-genome). Figure 2 shows that in both games, the evolved cooperation level is strongly affected by both the genetic and the behavioural architecture. In fact, in both games, average cooperation levels were 0.4 or lower for one implementation and 0.8 or higher for another implementation. Cooperation levels were higher for the 1:1 architecture when compared with the neural network architecture in all scenarios, but the effect of the mutation regime was different between the two games. In the IPD, per-locus mutation was associated with higher levels of cooperation than whole-genome mutation, whereas the opposite was true in the ISD. To understand the causes underlying these large differences, we next zoom in on the evolutionary dynamics of the 16 strategies that

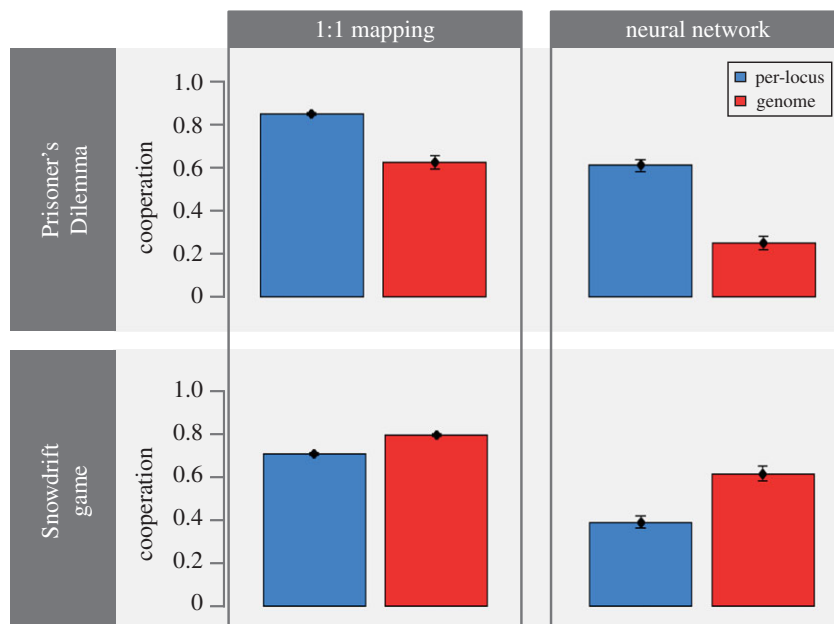


Figure 2. Cooperation levels in the IPD (top) and the ISD (bottom) for all four mechanistic implementations of the 16 strategies. The bars show average cooperation levels over all interactions in the last generation, across all replicates. Error bars show standard error of the mean.

were considered in this study (see table 1 for a complete list and an explanation of strategy names).

(b) Evolutionary dynamics in the Iterated Snowdrift game

A game-theoretical analysis of the 16 strategies in the ISD reveals that there are three ESSs (see the electronic supplementary material for details). ESS 1 consists of 83.3% *con-D* (conventional defector, a strategy that sticks with its previous behaviour if it played the opposite as its opponent in the previous round, and defects otherwise), together with 16.7% *ALLD*. ESS 2 consists of the pure strategy *Pavlov*. ESS 3 involves three pure strategies: 96.8% *MNG* (Mr Nice Guy, which always cooperates, except if it cooperated while the interaction partner defected in the previous round), 2.2% *inconsistent* (which always plays the opposite to its previous move) and 1.0% *Acon-D* (unconventional defector, a strategy that changes behaviour if it played the opposite as its opponent in the previous round, and defects otherwise).

In our simulations, we recover the three ESSs above. Typically, a simulation stays at one of the ESSs for extensive periods of time, followed by a rapid shift to another ESS. In most simulations across all scenarios, ESS 1 evolved first. In some of the simulations, ESS 1 was invaded by *Pavlov*, leading to the establishment of ESS 2. In a subset of those cases, ESS 2 was ultimately invaded by *MNG*, establishing ESS 3. ESS 3 was almost never invaded, but in very rare cases could be followed by a new establishment of ESS 1. The probability of transition between two ESSs and, accordingly, the probability to find the population in any of the three ESSs strongly depends on the behavioural architecture and mutation regime (figure 3): in case of a 1:1 architecture, ESS 3 is the dominant state in case of per-locus mutation, while the simulations switch between ESS 2 (attained 81% of the time) and 3 (19%) in case of whole-genome mutation. In case of a neural network architecture, the simulations

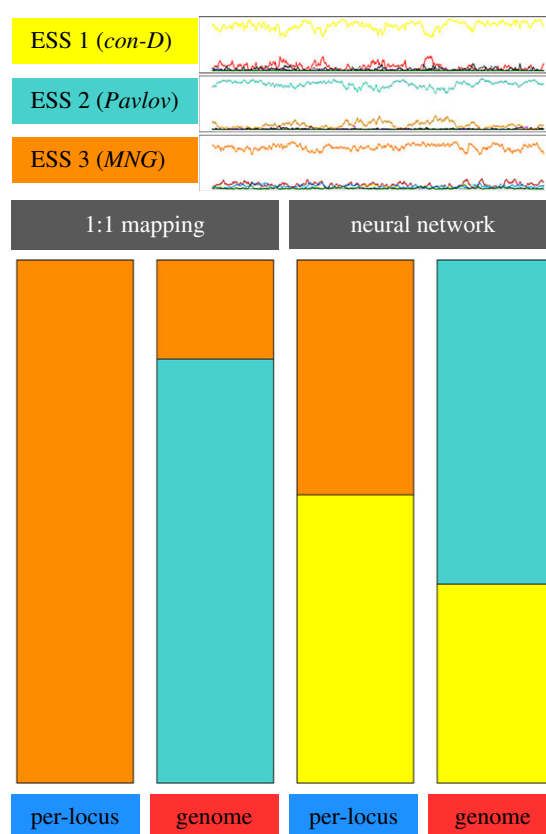


Figure 3. Simulation outcomes in the ISD, for all four mechanistic implementations. The three line graphs (top) show time series (2500 generations) of typical simulation runs, each illustrating the attainment of one of the three ESSs of this game (see the electronic supplementary material, table S1). The bar graphs (bottom) show for each scenario the fraction of 100 replicate simulations for which the last generation was in each of the three ESSs.

either end up in ESS 1 (55%) or 3 (45%) in case of per-locus mutation and in ESS 1 (38%) and 2 (62%) in case of whole-genome mutation. In other words, the four implementations differ in their likelihood of attaining each of the ESSs, and this

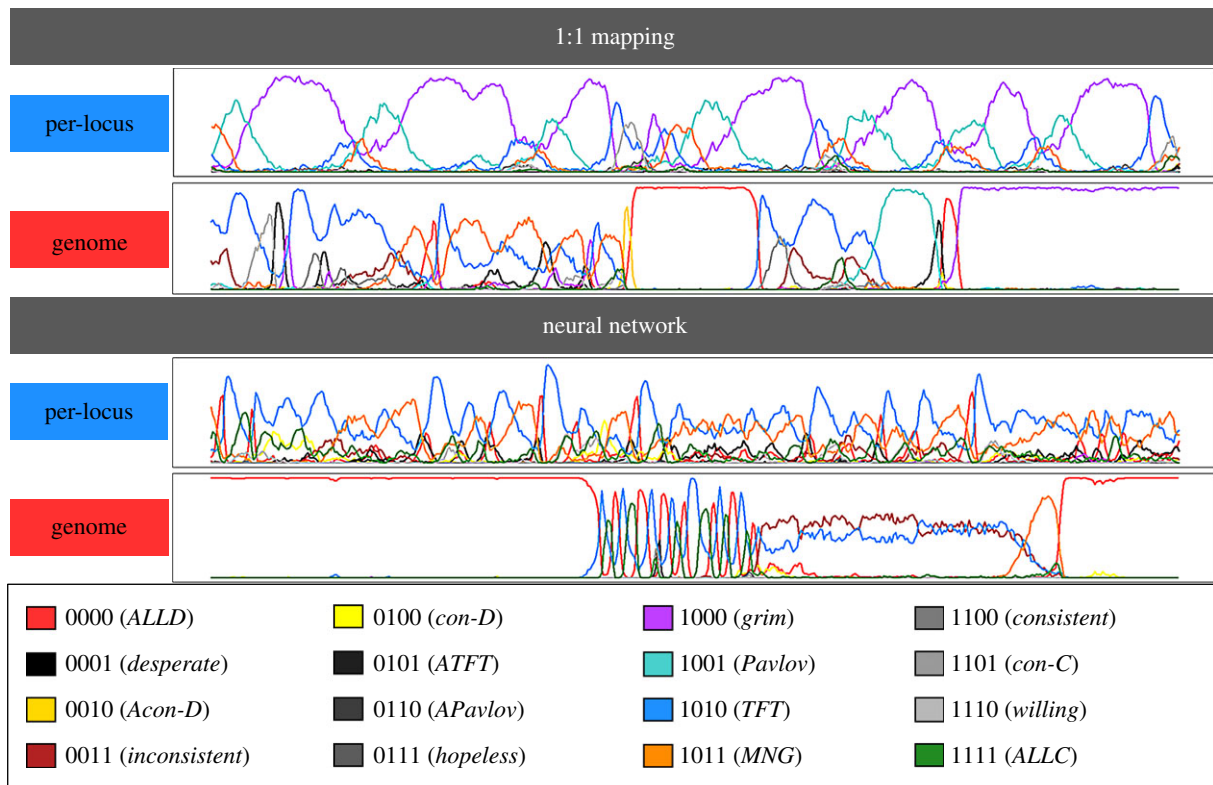


Figure 4. Time series of typical simulation runs in the IPD, for all four mechanistic implementations. In each case, a time period of 2500 generations is shown. The coloured lines represent the frequencies of the 16 different strategies.

difference is reflected in the average cooperation levels observed in figure 2 (as each ESS induces a different cooperation level; see electronic supplementary material, table S1).

Why do the behavioural and the genetic architecture have such a strong effect on the evolutionary outcome? This can be illustrated by considering the transition from ESS 1 to ESS 2. In ESS 1, *Pavlov* has a slight selective disadvantage when rare, but as soon as it occurs in higher frequencies, it achieves a higher pay-off than the strategies in ESS 1 (because the pay-off of *Pavlov* against itself is high). Therefore, if *Pavlov* increases enough against the selection gradient owing to mutation and genetic drift, it can invade, and ESS 2 becomes established. Clearly, the probability that *Pavlov* results from mutation of the strategies in ESS 1 is a crucial factor in this regard. As shown in table 1 (and explained in the electronic supplementary material), *Pavlov* occupies a much larger part of the genotype space (6.25%) in the 1:1 architecture than in the neural network architecture (0.08%). As a result, *Pavlov* almost never invades ESS 1 in the neural network architecture, whereas this often happens in the 1:1 architecture.

(c) Evolutionary dynamics in the Iterated Prisoner's Dilemma game

A game-theoretical analysis of the 16 strategies in the IPD reveals two ESSs, both containing only a single strategy: *ALLD* and *grim* (see the electronic supplementary material for details).

In our simulations, there indeed were extended periods of time in which either *ALLD* or *grim* are dominant in the population. However, in most cases, the evolution of strategies was

very dynamic and often irregular. This is in line with earlier studies which also conclude that the evolutionary dynamics in an IPD is often chaotic and off-equilibrium [48–50]. Like in the ISD, both the behavioural and the genetic architecture had a strong effect on the evolutionary dynamics (figure 4). In the case of 1:1 mapping with per-locus mutation, steady cycles of *grim*, *TFT*, *MNG* and *Pavlov* were observed for all replicate simulations (this is consistent with earlier findings by Nowak & Sigmund [48]). For whole-genome mutation, the patterns look less consistent (yet highly dynamic), including longer spells of *ALLD* domination (this explains the relatively low cooperation levels in this scenario). In the neural network architecture, per-locus mutation led to very dynamic yet fairly consistent patterns, mostly involving *TFT* and *MNG*, and infrequent *ALLD* domination spells. Entire-genome mutation typically led to long *ALLD* domination spells interspersed by short periods with both cyclical dynamics involving various strategies including *TFT*, *ALLC*, *ALLD*, *grim*, and *MNG* and non-cyclical coexistence of *TFT* and *inconsistent*.

The effect of underlying mechanisms on the evolutionary dynamics can be explained by the fact that different mechanisms induce differences in the 'mutational distance' between strategies, that is, the likelihood that a mutation in a strategy gives rise to a given alternative strategy. As an example, consider the extended periods of dominance of *grim* that were frequently observed. Those periods are typically ended by the invasion of *TFT*. *TFT* obtains a slightly worse pay-off against *grim* than *grim* obtains against itself. However, *TFT* does obtain better pay-offs when it happens to be paired with itself. In other words: if *TFT* can increase enough against the selection gradient because of genetic drift and mutation, it gains a selective advantage and can invade. The probability

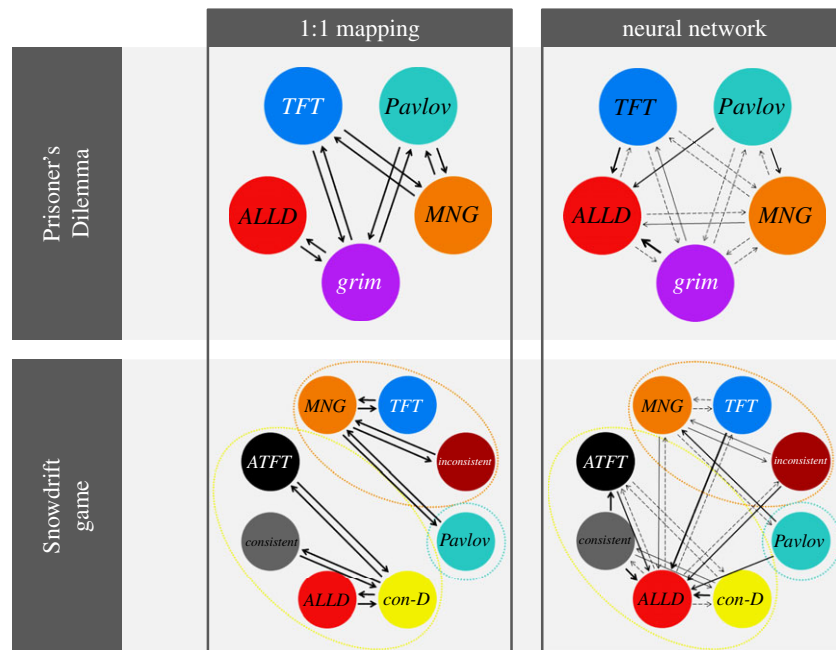


Figure 5. Mutational distance between the most relevant strategies in the Prisoner's Dilemma (top) and the ISD (bottom), for both behavioural architectures and the case of per-locus mutation. An arrow pointing from one strategy to another indicates that a mutation of the former strategy has probability of larger than 0.001 to yield the latter strategy. A probability of more than 0.05 is indicated by a solid arrow (the thickness of the arrow is proportional to the probability). In the 1 : 1 model, each strategy can mutate to four other strategies with equal probability, so the arrows in the mutation maps for the 1 : 1 model all represent a probability of 0.25. To calculate these probabilities, we first generated a large number of random genotypes (in the same way as generating a genotype through 'entire-genome mutation'), and determined their corresponding strategy. Then, for each strategy, we mutated all corresponding genotypes many times, and again determined the resulting strategies.

that this occurs depends on the implementation: in the 1 : 1 architecture with per-locus mutation, a mutation of *grim* produces *TFT* with probability 1/4. In the case of whole-genome mutation, this probability is only 1/16—this makes it considerably less likely that *TFT* obtains appreciable frequencies, and explains the extended spells of *grim* domination (figure 4) and the lower degree of cooperation (figure 2) in this case.

4. Discussion

Our study demonstrates that behavioural architecture and mutation regime are of considerable importance for the dynamics and outcome of social evolution. In the IPD, we observed three types of dynamic behaviour (predictable cycles; fast and chaotic dynamics; spells of *ALLD* or *grim* domination) whose occurrence crucially depended on both behavioural architecture and mutation structure. Likewise, the prevalence of and the transitions between the three ESSs in the ISD were strongly determined by both architecture and mutation structure. In both games, the differences in evolutionary dynamics resulted in substantial differences in the average level of cooperation. These conclusions are not specific to the parameters considered in our simulations; they also hold for different pay-off configurations of both games, and for a lower degree of stochasticity in the simulations (see the electronic supplementary material).

The effect of mechanisms on the evolutionary dynamics was not caused by 'hard' constraints (the inability of mechanisms to produce all phenotypes), as all 16 strategies of the game were feasible in all four implementations. Yet, the mechanisms induced some 'soft' constraints on evolution,

by strongly affecting the probabilities with which strategies arise by mutation (see figure 5 for a schematic overview of mutation probabilities for each scenario considered in this study). Even in case of small mutation rates, the mutational distribution has a strong effect on the type of variation that can be expected to be present in a given situation. Some strategies only gain a selective advantage once they have increased beyond a certain frequency, and mutation probabilities determine the probability that this will happen. In the 1 : 1 architecture considered in this study, each strategy has an equal probability to result from a randomly generated genotype, whereas in the neural network architecture, some strategies (notably *ALLD* and *ALLC*) are much more likely to arise owing to mutation than others (table 1). In the case of entire-genome mutation, the strategy of a mutant individual is independent of the strategy of its parent, and mutation probabilities therefore only depend on the behavioural architecture. In the case of per-locus mutation, the parental strategy partly determines the strategy of their mutant offspring.

We are not the first to point out that the genotype–phenotype mapping and the induced mutation structure are important for the course of evolution. In fact, the formal frameworks for modelling evolutionary dynamics can, to a certain extent, take these complexities into account. For example, the 'canonical equation' of adaptive dynamics theory includes a mutational covariance matrix [51,52], which characterizes the likelihood that a combination of phenotypic traits (like a conditional strategy) arises and potentially invades the current resident strategy. Likewise, the multivariate selection equation of quantitative genetics [53,54] can be written in a form that it includes a matrix characterizing the covariance in phenotypic traits between parents and offspring [55]. Sean Rice has worked out formally

how these covariances arise and determine the course of evolution when phenotypic traits are the outcome of (developmental) mechanisms [55,56]. We are not aware of attempts to actually derive the covariance matrices of adaptive dynamics of quantitative genetics theory on the basis of a concrete mechanistic model. Instead, theoretical studies tend to make simplifying assumptions, such as replacing the covariance matrix by the identity matrix (e.g. [25]). Already in the case of frequency-independent selection and in the absence of stochasticity, assumptions like these are not unimportant, as the covariance structure largely determines which peak of a multi-peaked fitness landscape will be reached.

There are two main reasons why we think that phenotypic covariances and, hence, the mechanisms underlying the development of phenotypes are of particular importance for social evolution. First, selection will virtually always be frequency-dependent in this case. As a consequence, the success of each strategy will strongly depend on the context, and, in particular, on the presence of selectively favoured competitors. Accordingly, a given architecture will contribute to the stability of a given equilibrium if it makes the production of selectively favoured alternatives less likely, and it will have a destabilizing effect if the opposite is the case. Second, in case of social interactions, there are typically many alternative Nash equilibria and ESSs. This is already illustrated in the IPD and the ISD with highly restricted memory considered here. Relaxing the restrictions on the strategy set would lead to a rapid increase in the number of equilibria. In fact, the Folk theorem of game theory [23,57] implies quite generally that in repeated games the set of Nash equilibrium strategies is so large that virtually any 'reasonable' outcome (in case of our IPD: any outcome between 0 and 5; in case of our ISD: any outcome between 1 and 5) can be achieved as the average outcome of an equilibrium. But also non-repeated games typically have several (and often a large number of) Nash equilibria and ESSs [58]. In all these cases, it is to be expected that the evolutionary dynamics will be affected in a similar way by mechanisms as in this study.

We have here focused on situations where the evolutionary game dynamics [42] are relatively simple. For the pay-off structure considered, the IPD and the ISG have a small number of ESSs, and these are the only attractors of the replicator equation (see the electronic supplementary material). Accordingly, mechanisms will mainly affect the transition between ESSs, as described above. It is conceivable that mechanisms have an even stronger effect in the presence of limit cycles or other non-equilibrium attractors. Such attractors regularly occur in evolutionary games (such as variants

of the Rock–Scissors–Paper game; [59]), and it has been shown that seemingly small differences in the genetic implementation of strategies can have major effects on the evolutionary outcome [60]. In the electronic supplementary material, we show that non-equilibrium attractors can also occur in the IPD (for slightly different pay-off parameters), but a thorough investigation of the interplay of genetic or behavioural architecture and non-equilibrium dynamics is beyond the scope of this study.

Our results should be mainly viewed as proof of principle that mechanisms matter for the course and outcome of social evolution. It would be premature to conclude that one of the four implementations considered in our study is more 'realistic' than the others. On purpose, we kept our assumptions on architecture as simple as possible, as this allowed us to develop a sound intuitive understanding of our results (figure 5). Because of this understanding, we are confident that our findings are of general relevance. The development of truly 'realistic' models remains a major challenge, as the actual genetic, physiological, neurological and psychological mechanisms behind social behaviour are still largely *terra incognita* for virtually all organisms and virtually all interaction types. For this reason, it would be premature to abandon the standard 1 : 1 genotype–phenotype mapping assumption in favour of (for example) a neural network implementation. However, whatever the implementation chosen, researchers should be aware that it may have considerable implications for the course and outcome of evolution.

The evolution of social behaviour is often an intricate process, with many feedbacks at work, and many possible outcomes. We have shown that underlying mechanisms are of decisive importance in determining which outcome eventually emerges in evolution. Therefore, it is of importance that we focus more on mechanisms when trying to explain the evolution of social behaviour. Both empirical work focused on understanding mechanisms and theoretical work investigating their importance for the dynamics and course of evolution have a vital role to play in this regard.

Data accessibility. All data are available on Dryad: <http://dx.doi.org/10.5061/dryad.5hq6r>.

Authors' contributions. P.v.d.B. and F.J.W. designed the study, conducted analyses, drafted the article and gave final approval of the paper. P.v.d.B. conceived of the study and conducted simulations.

Competing interests. We declare we have no competing interests.

Funding. P.v.d.B. was supported by grant no. NWO 022.033.49 of the Graduate Programme of the Netherlands Organisation for Scientific Research.

References

1. Mayr E. 1961 Cause and effect in biology. *Science* **134**, 1501–1506. (doi:10.1126/science.134.3489.1501)
2. Tinbergen N. 1963 On aims and methods of ethology. *Z. Tierpsychol.* **20**, 410–433. (doi:10.1111/j.1439-0310.1963.tb01161.x)
3. Mayr E. 1959 Where are we. *Cold Spring Harb. Symp. Quant. Biol.* **24**, 1–14. (doi:10.1101/SQB.1959.024.01.003)
4. Maynard Smith J, Burian R, Kauffman S, Alberch P, Campbell J, Goodwin B, Lande R, Raup D, Wolpert L. 1985 Developmental constraints and evolution. *Q. Rev. Biol.* **60**, 265–287. (doi:10.1086/414425)
5. Maynard Smith J, Szathmáry E. 1995 *The major transitions in evolution*. Oxford, UK: Freeman.
6. Arthur W. 2002 The emerging conceptual framework of evolutionary developmental biology. *Nature* **415**, 757–764. (doi:10.1038/415757a)
7. Müller GB. 2007 Evo-devo: extending the evolutionary synthesis. *Nat. Rev. Genet.* **8**, 943–949. (doi:10.1038/nrg2219)
8. Aldana M, Balleza E, Kauffman S, Resendiz O. 2007 Robustness and evolvability in genetic regulatory

- networks. *J. Theor. Biol.* **245**, 433–448. (doi:10.1016/j.jtbi.2006.10.027)
9. Ciliberti S, Martin OC, Wagner A. 2007 Robustness can evolve gradually in complex regulatory gene networks with varying topology. *PLoS Comput. Biol.* **3**, e15. (doi:10.1371/journal.pcbi.0030015)
 10. Lozada-Chavez I, Janga SC, Collado-Vides J. 2006 Bacterial regulatory networks are extremely flexible in evolution. *Nucleic Acids Res.* **34**, 3434–3445. (doi:10.1093/nar/gkl423)
 11. Pfennig DW, Ehrenreich IM. 2014 Towards a gene regulatory network perspective on phenotypic plasticity, genetic accommodation and genetic assimilation. *Mol. Ecol.* **23**, 4438–4440. (doi:10.1111/mec.12887)
 12. Schneider RF, Li Y, Meyer A, Gunter HM. 2014 Regulatory gene networks that shape the development of adaptive phenotypic plasticity in a cichlid fish. *Mol. Ecol.* **23**, 4511–4526. (doi:10.1111/mec.12851)
 13. Botero CA, Weissing FJ, Wright J, Rubenstein DR. 2015 Evolutionary tipping points in the capacity to adapt to environmental change. *Proc. Natl Acad. Sci. USA* **112**, 184–189. (doi:10.1073/pnas.1408589111)
 14. McNamara JM, Houston AI. 2009 Integrating function and mechanism. *Trends Ecol. Evol.* **24**, 670–675. (doi:10.1016/j.tree.2009.05.011)
 15. Fawcett TW, Hamblin S, Giraldeau L-A. 2013 Exposing the behavioral gambit: the evolution of learning and decision rules. *Behav. Ecol.* **24**, 2–11. (doi:10.1093/beheco/ars085)
 16. Fawcett TW, Fallenstein B, Higginson AD, Houston AI, Mallpress DEW, Trimmer PC, McNamara JM. 2014 The evolution of decision rules in complex environments. *Trends Cogn. Sci.* **18**, 153–161. (doi:10.1016/j.tics.2013.12.012)
 17. Enquist M, Arak A. 1994 Symmetry, beauty, and evolution. *Nature* **372**, 169–172. (doi:10.1038/372169a0)
 18. Gross R, Houston AI, Collins EJ, McNamara JM, Dechaume-Montcharmont F-X, Franks NR. 2008 Simple learning rules to cope with changing environments. *J. R. Soc. Interface* **5**, 1193–1202. (doi:10.1098/rsif.2007.1348)
 19. McNally L, Brown SP, Jackson AL. 2012 Cooperation and the evolution of intelligence. *Proc. R. Soc. B* **279**, 3027–3034. (doi:10.1098/rspb.2012.0206)
 20. McNamara JM, Gasson CE, Houston AI. 1999 Incorporating rules for responding into evolutionary games. *Nature* **401**, 368–371. (doi:10.1038/43869)
 21. Taylor PD, Day T. 2004 Stability in negotiation games and the emergence of cooperation. *Proc. R. Soc. Lond. B* **271**, 669–674. (doi:10.1098/rspb.2003.2636)
 22. Akçay E, Van Cleve J, Feldman MW, Roughgarden J. 2009 A theory for the evolution of other-regard integrating proximate and ultimate perspectives. *Proc. Natl Acad. Sci. USA* **106**, 19 061–19 066. (doi:10.1073/pnas.0904357106)
 23. Van Damme E. 1991 *Stability and perfection of Nash equilibria*. Berlin, Germany: Springer.
 24. Nowak MA, Sigmund K. 1993 A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **364**, 56–58. (doi:10.1038/364056a0)
 25. Van Doorn GS, Hengeveld GM, Weissing FJ. 2003 The evolution of social dominance. I. Two-player models. *Behaviour* **140**, 1305–1332. (doi:10.1163/156853903771980602)
 26. Harsanyi JC, Selten R. 1988 *A general theory of equilibrium selection in games*. Cambridge, MA: MIT Press.
 27. Samuelson L. 1998 *Evolutionary games and equilibrium selection*. Cambridge, MA: MIT Press.
 28. Gintis H. 2000 *Game theory evolving*. Princeton, NJ: Princeton University Press.
 29. Axelrod R, Hamilton WD. 1981 The evolution of cooperation. *Science* **211**, 1390–1396. (doi:10.1126/science.7466396)
 30. Boyd R, Lorberbaum JP. 1987 No pure strategy is evolutionarily stable in the repeated Prisoner's Dilemma game. *Nature* **327**, 58–59. (doi:10.1038/327058a0)
 31. Binmore K, Samuelson L. 1992 Evolutionary stability in repeated games played by finite automata. *J. Econ. Theory* **57**, 278–305. (doi:10.1016/0022-0531(92)90037-1)
 32. Nowak MA, Sigmund K. 1998 Evolution of indirect reciprocity by image scoring. *Nature* **393**, 573–577. (doi:10.1038/31225)
 33. Doebeli M, Hauert C. 2005 Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game. *Ecol. Lett.* **8**, 748–766. (doi:10.1111/j.1461-0248.2005.00773.x)
 34. McNamara JM, Barta Z, Fromhage L, Houston AI. 2008 The coevolution of choosiness and cooperation. *Nature* **451**, 189–192. (doi:10.1038/nature06455)
 35. Maynard Smith J, Price GR. 1973 The logic of animal conflict. *Nature* **246**, 15–18. (doi:10.1038/246015a0)
 36. Sugden R. 1986 *The economics of rights, cooperation and welfare*. Oxford, UK: Basil Blackwell.
 37. Santos FC, Pacheco JM. 2005 Scale-free networks provide a unifying framework for the emergence of cooperation. *Phys. Rev. Lett.* **95**, 098104. (doi:10.1103/PhysRevLett.95.098104)
 38. Maynard Smith J. 1982 *Evolution and the theory of games*. Cambridge, UK: Cambridge University Press.
 39. Dugatkin L. 1998 *Game theory and animal behaviour*. Oxford, UK: Oxford University Press.
 40. Weibull JW. 1995 *Evolutionary game theory*. Cambridge, MA: MIT Press.
 41. Hofbauer J, Sigmund K. 1998 *Evolutionary games and population dynamics*. Cambridge, UK: Cambridge University Press.
 42. Hofbauer J, Sigmund K. 2003 Evolutionary game dynamics. *Bull. Am. Math. Soc.* **40**, 479–519. (doi:10.1090/S0273-0979-03-00988-1)
 43. Nowak MA. 2006 *Evolutionary dynamics: exploring the equations of life*. Cambridge, MA: Harvard University Press.
 44. Nowak MA, Sasaki A, Taylor C, Fudenberg D. 2004 Emergence of cooperation and evolutionary stability in finite populations. *Nature* **428**, 646–650. (doi:10.1038/nature02414)
 45. Imhof LA, Fudenberg D, Nowak MA. 2005 Evolutionary cycles of cooperation and defection. *Proc. Natl Acad. Sci. USA* **102**, 10 797–10 800. (doi:10.1073/pnas.0502589102)
 46. Weissing FJ. 1996 Genetic versus phenotypic models of selection: can genetics be neglected in a long-term perspective? *J. Math. Biol.* **34**, 533–555. (doi:10.1007/BF02409749)
 47. Enquist M, Ghirlanda S. 2005 *Neural networks and animal behavior*. Princeton, NJ: Princeton University Press.
 48. Nowak MA, Sigmund K. 1993 Chaos and the evolution of cooperation. *Proc. Natl Acad. Sci. USA* **90**, 5091–5094. (doi:10.1073/pnas.90.11.5091)
 49. Gosak M, Marhl M, Perc M. 2008 Chaos between stochasticity and periodicity in the Prisoner's Dilemma game. *Int. J. Bifurcat. Chaos* **3**, 869–875. (doi:10.1142/S0218127408020720)
 50. Rasmusen E. 1989 *Games and information*. New York, NY: Basil Blackwell.
 51. Dieckmann U, Law R. 1996 The dynamical theory of coevolution: a derivation from stochastic ecological processes. *J. Math. Biol.* **34**, 579–612. (doi:10.1007/BF02409751)
 52. Metz JAJ, de Kovel CGF. 2013 The canonical equation of adaptive dynamics for Mendelian diploids and haplo-diploids. *Interface Focus* **3**, 20130025. (doi:10.1098/rsfs.2013.0025)
 53. Lande R. 1979 Quantitative genetic analysis of multivariate evolution, applied to brain-body size allometry. *Evolution* **33**, 402–416. (doi:10.2307/2407630)
 54. Lande R, Arnold SJ. 1983 The measurement of selection on correlated characters. *Evolution* **37**, 1210–1226. (doi:10.2307/2408842)
 55. Rice SH. 2004 *Evolutionary theory*. Sunderland, MA: Sinauer Associates.
 56. Rice SH. 2002 A general population genetic theory for the evolution of developmental interactions. *Proc. Natl Acad. Sci. USA* **99**, 15 518–15 523. (doi:10.1073/pnas.202620999)
 57. Brems B. 1996 Cheating and cooperation: potential solutions to the Prisoner's Dilemma. *Oikos* **76**, 14–24. (doi:10.2307/3545744)
 58. Selten R. 1983 Evolutionary stability in extensive two-person games. *Math. Soc. Sci.* **5**, 269–363. (doi:10.1016/0165-4896(83)90012-4)
 59. Weissing FJ. 1991 Evolutionary and dynamic stability in a class of evolutionary normal form games. In *Game equilibrium models. I. Evolution and game dynamics* (ed. R Selten), pp. 29–97. Berlin, Germany: Springer.
 60. Weissing FJ, Van Boven M. 2001 Selection and segregation distortion in a sex-differentiated population. *Theor. Pop. Biol.* **60**, 327–341. (doi:10.1006/tpbi.2001.1550)

SI APPENDIX

The importance of mechanisms for the evolution of cooperation

Pieter van den Berg & Franz J. Weissing

CONTENTS

Implementation of strategies by a neural network	2
ESS analysis of the IPD and the ISD	4
Sensitivity analysis	7
References	15

1. Implementation of strategies by a neural network

All sixteen strategies listed in Table 1 can be realized by our neural network architecture. To give some more insight in how a neural network configuration implements a strategy, Figure S1 gives an example of one of the many possible networks implementing *Pavlov*. *Pavlov* cooperates after mutual cooperation or mutual defection, and defects otherwise. The implementation of *Pavlov* in the neural network architecture depends on a relatively precise relationship between the weights and thresholds of the network. For example, the sum of w_5 and w_6 must exceed t_3 , but neither of the two should exceed t_3 alone. Also, w_1 and w_3 must together exceed t_1 , and only one of those two weights must exceed t_1 alone, while the other must not. The same necessary relationships hold for weights w_2 and w_4 and threshold t_2 . In addition, both t_1 and t_2 must be negative. The relatively restrictive conditions needed for a network to correspond to *Pavlov* give an intuitive understanding of why only a small percentage of the genotype space is associated with this strategy (see Table 1).

When considering the network in Figure S1, it also becomes easier to understand why the strategies *ALLD* and *ALLC* are so likely to result from a random configuration of the network. If the value of t_3 in this network is changed from 4 to 6, then this threshold can never be exceeded, meaning this network will always defect. More generally, all networks for which $t_3 > w_5 + w_6$ and $t_3 > 0$ implement *ALLD*, regardless of the values of the other weights and thresholds. Similarly, all networks for which $t_3 < w_5 + w_6$ and $t_3 < 0$ implement *ALLC*. However, this is only one of many ways that the network in Figure S1 could mutate towards a network implementing *ALLD*. For example, since both t_1 and t_2 must be exceeded for this network to cooperate, any mutation that increases the value of either of these two thresholds so that it cannot be exceeded by the preceding weights would result in an *ALLD* network.




Figure S1. A neural network corresponding to the *Pavlov* strategy. The network on top shows the values of all the weights ($w_1 - w_6$) and thresholds ($t_1 - t_3$); green lines depict positive values, red lines depict negative values, and their thickness indicate their absolute values. The input nodes receive the previous own decision (top) and the previous decision of the interaction partner (bottom); 1 for cooperation and 0 for defection. The four networks on the bottom show in detail how the output of the network is generated for each of the four possible outcomes of the previous round. The values of the input nodes are multiplied with weights w_1 to w_4 ; the resulting values are summed and fed into the hidden layer. If the summed values exceed the respective threshold values of the hidden layer, their output is 1; otherwise, it is 0. Those values are then multiplied with weights w_5 and w_6 , and the sum of the resulting values is fed into the output node. If this value exceeds the threshold of the output node, the network cooperates; otherwise, it defects.

2. ESS Analysis of the IPD and the ISD

Repeated games like the Iterated Prisoners' Dilemma game (IPD) or the Iterated Snowdrift Game (ISD) have a huge strategy set, and even for games as simple as these a full game theoretical analysis has not yet been achieved. Such an analysis is a formidable task, since the number of Nash equilibria is huge [a]. In fact, the 'Folk Theorem' of game theory implies that *any* reasonable outcome can be realized by a Nash equilibrium of the iterated game (for details see [b]). On the other hand, no pure strategy is evolutionarily stable in the IPD or the ISD (e.g. [c]).

Here, we consider simpler versions of the IPD and the ISD where the strategy set is restricted to 16 pure strategies with limited memory (only the moves in the previous round are memorized). These versions of the IPD and the ISD have been the subject of many studies, but to our best knowledge a full characterization of all Nash equilibria or all evolutionarily stable strategies (ESS) has never been given. This is understandable, since even under the restriction to 16 pure strategies the game theoretical analysis remains a challenge. Here we show how to classify all ESSs with support 1, 2 and 3 of the two repeated games for the payoff configuration considered in the main text. A Mathematica file with all calculations is available upon request. An overview of all ESSs with support 1, 2 and 3 is shown in Table S1. By means of numerical iterations based on the replicator equation, we demonstrate that most likely there are no other ESSs with a larger support.

Construction of the 16x16 payoff matrix

As a first step, we determine the payoff matrix of the game by calculating the expected payoffs for any pair of pure strategies. Any such pair of strategies induces a sequence of transitions between the four different 'cooperation states' of a game round (mutual cooperation, mutual defection, cooperation-defection, defection-cooperation). In the absence of errors, this sequence would be deterministic and mainly dependent on the combination of moves in the first round. The inclusion of perception and implementation errors, however, turns the transition between states into a stochastic process with a well-defined 4x4-matrix of transition probabilities between cooperation states. The right eigenvector corresponding to the dominant eigenvalue of this matrix (which is typically the only positive eigenvalue of the matrix) is proportional to the stationary distribution over states generated by the interaction of the two pure strategies. The four elements of the (normalized) eigenvector correspond to the fraction of time spent in each of the four cooperation states in an infinitely repeated interaction. Weighing the four elements of the payoff matrix of the one-shot game by these elements and summing up the results therefore yields the expected per-round payoff for each of the two strategies. Notice that this pair of payoffs does not depend on the initial pair of moves (these are irrelevant from a long-term perspective),

but reflects the perception and implementation errors made by the players (both were kept at 0.01, as in the simulations). All subsequent ESS calculations are based on the 16x16 payoff matrix that results by applying the above recipe to all pairs of pure strategies.

Determination of all pure-strategy ESSs

It is straightforward to characterize all pure-strategy Nash equilibria: A pure strategy i is a Nash equilibrium if no alternative pure strategy j attains a higher payoff against i than i attains against itself. If i is a ‘strong’ Nash equilibrium (any alternative pure strategy j attains a lower payoff against i than i attains against itself), then i is an ESS [b, d-g]. It turned out that the IPD has two pure-strategy Nash equilibria (*grim* and *ALLD*), which both are strong Nash equilibria and therefore an ESS. The ISD has a single pure-strategy Nash equilibrium (*Pavlov*), which again is a strong Nash equilibrium and an ESS.

Determination of all ESSs with two coexisting pure strategies

To calculate all ESSs with support two, we first determined all those pairs of pure strategies i and j that can mutually invade each other: j has a higher payoff against i than i has against itself; and i has a higher payoff against j than j has against itself. This condition of mutual invadability is equivalent to the existence of a mixed-strategy ESS of the restricted game with only these two pure strategies [f,g]. This ESS can easily be calculated on the basis of the condition that the fitness of both pure strategies needs to be identical at the ESS [d]. The ESS thus found is an ESS of the full game (with all 16 strategies) if all other strategies have a lower payoff at this two-strategy equilibrium than the two equilibrium strategies [h]. It turned out that the IPD has no ESS with support two, while the ISD has one such ESS: 83.3% *con-D* and 16.7% *ALLD*.

Determination of all ESSs with three coexisting pure strategies

The calculation of all ESSs with support three is more tedious. For each triplet i, j , and k of pure strategies we first checked whether the three strategies can coexist at equilibrium. This was done by checking whether the system of linear equations specifying fitness equality of the three strategies in the restricted three-strategy game has a positive solution. This solution specifies a candidate-ESS. In a second step, the ‘internal’ stability of this candidate-ESS (i.e., its evolutionary stability in the restricted three-strategy game) was checked, making use of a criterion for evolutionary stability in 3x3 games [f,g]. In a third step, the ‘external’ stability of the candidate-ESS was determined by checking whether all other 13 pure strategies of the full 16-strategy game have a lower payoff at the candidate-ESS than the three strategies being part of the candidate-ESS [h]. It turned out that the IPD has no ESS with support three, while the ISD has one such ESS: 96.8% *MNG*, 2.2% *inconsistent*, and 1.0% *Acon-D*.

Table S1. The evolutionarily stable strategies (ESS) identified by the game theoretical analysis of the IPD and the ISD. In case of the ISD, the three ESSs are numbered in line with the three equilibrium outcomes of the simulations described in the main text. The last column gives the average cooperation level in each equilibrium. A full description of all the strategies can be found in Table 1 of the main text.

<i>Game</i>	<i>ESS</i>	<i>Strategies</i>	<i>Fraction</i>	<i>Cooperation</i>
IPD	1	<i>ALLD</i>	1.000	0.010
	2	<i>grim</i>	1.000	0.013
ISD	1	<i>con-D</i>	0.833	0.180
		<i>ALLD</i>	0.167	
	2	<i>Pavlov</i>	1.000	0.943
	3	<i>MNG</i>	0.968	0.698
		<i>inconsistent</i>	0.022	
		<i>Acon-D</i>	0.010	

Numerical iterations based on the replicator equation

To check whether the equilibria identified in the game theoretical analysis are actually attainable and dynamically stable, and whether any other attractors are present in the system, we performed extensive numerical iterations using the replicator equation [f,g]. To do this, we used the same 16x16 payoff matrices calculated for the game theoretical analysis of both games (described above). Starting from around $6 \cdot 10^7$ different initial population constitutions, we iterated the replicator equation until an attractor was reached. For each iteration, the minimum frequency of each strategy was set to 0.001 so that invasion of any strategy was always in principle possible. The outcome of these iterations was congruent with the game theoretical analysis; all iterations (for both of the games) ended up in one of the ESSs in Table S1, and each of the ESSs was commonly attained (depending on the initial conditions). The oscillating behaviour commonly observed in the individual-based simulations of the IPD was never observed in the (deterministic) numerical iterations. This suggests that those observations correspond to transient (‘away from equilibrium’) behaviour driven by the stochastic component of the individual-based simulations.

3. Sensitivity analysis

In the main text, we have shown that the mechanisms underlying the strategies of an evolutionary game can be of substantial importance for the evolutionary dynamics and the evolutionary outcome. For ease of representation, all simulations were conducted for one set of parameters. In this section, we show that our main conclusion also holds for a number of altered parameter settings. Specifically, we consider versions of our model with a reduced mutation rate and with an altered payoff configuration in both games.

Reduced mutation rate

The simulations of the IPD discussed in the main text commonly exhibited highly dynamical behaviour, even though no non-equilibrium attractors were identified in the game theoretical analysis or numerical iterations of the replicator equation (see SI section 2). It is likely that this discrepancy is caused by the stochastic component of the simulations. The degree of stochasticity can be reduced by increasing the population size or by decreasing the mutation rate. Here, we investigate how our evolutionary simulations are affected by a reduction of the mutation rate from 10^{-3} to 10^{-4} . In addition, we also give an idea of how replicate simulation runs can differ from each other.

Figure S2 shows typical simulation runs of the IPD for each of the scenarios (behavioural architecture and mutation regime), for both mutation rates. The game theoretical analysis identified two pure-strategy ESSs for this game: *ALLD* and *grim* (see Table S1). Most simulation runs consist of periods of stasis (with one or two strategies in an equilibrium-like situation), followed by periods of strong fluctuations. As a rule, the periods of stasis are longer in case of a lower mutation rate, but even in that case, highly dynamical periods are common. During the periods of stasis, one of the two ESSs (*ALLD* or *grim*) is typically the predominant strategy. However, many simulation runs included prolonged periods of stasis dominated by the non-ESS strategies *Pavlov* or *MNG*, or *con-D* or static periods where *TFT* and *inconsistent* coexisted in almost equal frequencies. Interestingly, these deviations from the game theoretical ESS predictions were mainly observed in simulations with a low mutation rate.

Besides these effects of a lowered mutation rate on the simulation dynamics, Figure S2 clearly shows that lowering the mutation rate does not change our main conclusion that underlying mechanisms strongly affect the evolutionary outcome. For example, *ALLD* domination was more commonly observed in the simulations with the neural network implementation, whereas *grim* domination was more common for the 1:1 mapping.




Figure S2. Typical simulation runs for the IPD, for different behavioural architectures, mutation regimes, and mutation rates. For each parameter combination, excerpts of five replicate simulation runs are shown. Each excerpt comprises a time period of 2,000 generations. The excerpts were chosen from 100 replicate simulation runs (comprising 100,000 generations each) to give an impression of the ‘typical’ dynamics observed for each parameter combination. Graphs of all 100 replicate simulation runs for each parameter combination are available upon request.

The game theoretical analysis of the ISD identified three ESSs (see Table S1). Irrespective of the mutation rate, most simulations quickly converged to one of these ESSs; transitions between ESSs occurred, but only on rare occasions. When transitions occurred, ESS 1 was generally attained first, eventually succeeded by ESS 2, and finally followed by ESS 3 in some simulations. As expected, such transitions were less frequent in case of a lower mutation rate.

Figure S3 shows how frequently each of the three ESSs was attained after 100,000 generations in 100 replicate simulations of the ISD, for each behavioural architecture, mutation regime, and mutation rate. In all cases, ESS 1 was observed more often for low mutation rate than for high mutation rate (except for the 1:1 mapping with entire genome mutation, where this ESS was never the outcome, regardless of the mutation rate). Similarly, ESS 3 was always observed more often for the higher mutation rate (except for the neural network with entire genome mutation, where this equilibrium was never observed). Besides these effects, lowering the mutation rate does not alter our main conclusion: also in the ISD the underlying mechanisms strongly affect the outcome of evolution.

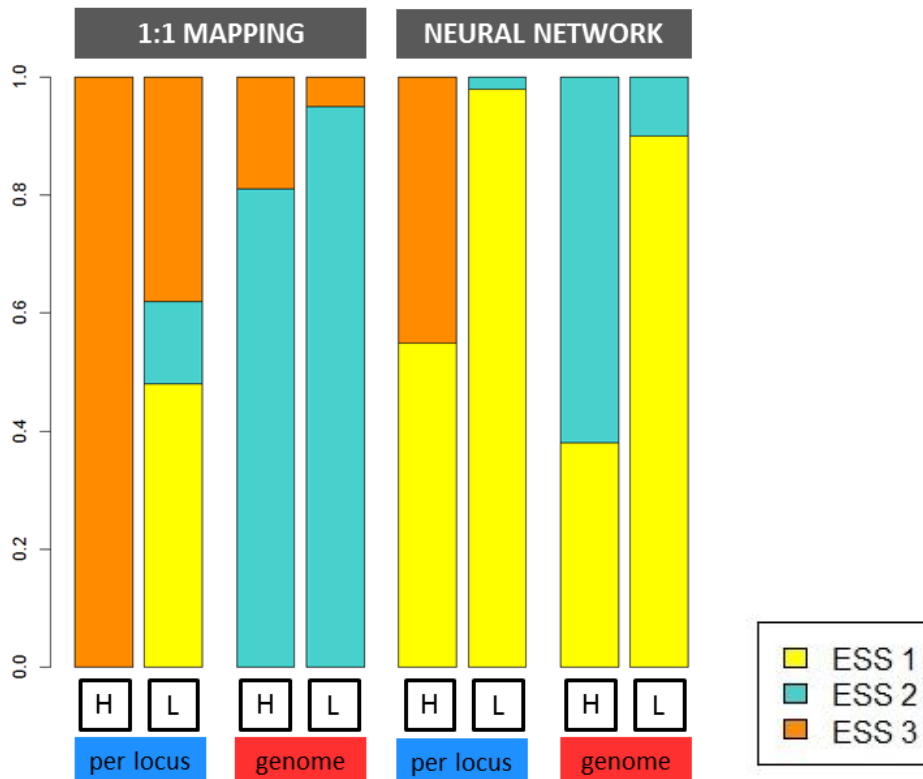


Figure S3. Evolutionary outcome of the simulations of the ISD for the different mechanistic implementations and mutation rates. The left bars of each pair show high mutation rates (H, $\mu = 0.001$), bars on the right show low mutation rates (L, $\mu = 0.0001$). Bars indicate the fraction of 100 replicate simulations ending up in each of the three ESSs of the ISD (Table S1) after 100,000 generations.

Alternative payoff configurations

A full sensitivity analysis concerning the effect of the payoff parameters of the IPD or the ISD on the evolutionary outcome is beyond the scope of this study. As mentioned above, even an ESS analysis or an analysis of the attractors of the replicator equation as a function of the payoffs is quite a challenge. This is exemplified by the game theoretical analysis in section 2 of this SI, which shows that the 16x16 payoff matrix of the two games does not have a straightforward relationship with the payoffs of the underlying one-shot game. Instead of conducting a comprehensive analysis of the effect of payoffs, we here just give an illustration, by redoing our analysis for one alternative payoff configuration for each of the games. As will become clear, even this relatively small alteration qualitatively changes the equilibrium structure of both games, but does not affect our main conclusion that underlying mechanisms strongly affect the outcome of social evolution.

In this analysis, we replaced the payoffs for defecting while the interaction partner cooperates from 5 to 4 in both games, yielding the following payoff matrices:

$$\text{PD: } \begin{pmatrix} 3 & 0 \\ 4 & 1 \end{pmatrix}; \quad \text{SD: } \begin{pmatrix} 3 & 1 \\ 4 & 0 \end{pmatrix}.$$

First, we performed a game theoretical ESS analysis on the iterated versions of these games (as described in section 2 of this SI); all ESSs with support 1, 2 or 3 are shown in Table S2. In case of the IPD, the game theoretical analysis identifies the same pure-strategy ESSs as in the original game (*ALLD* and *grim*), but also identifies a third (*Pavlov*). As in the original game, there are no ESSs with 2 or 3 coexisting strategies in this game. In the ISD, we also find the same pure-strategy ESS as in the original game (*Pavlov*). In addition, we find another pure-strategy ESS: *con-D*. This ESS is relatively similar to ESS 1 of the original model (that equilibrium also consisted of mostly *con-D*, but also included a relatively low frequency of *ALLD*; see Table S1). These two pure-strategy ESSs are the only ESSs identified by the game theoretical analysis; ESS 3 (or a similar ESS) was not identified in this version of the ISD.

To check whether there are any other attractors in the system, we also did extensive numerical iterations of the replicator equation (see section 2 of this SI for details). It turned out that all the ESSs described above were commonly attained; and that there is apparently no alternative stable equilibrium.

Interestingly, a non-equilibrium attractor appeared in the replicator dynamics of the IPD: about 18% of all iterations converged to the stable limit cycle shown in Fig. S4. This cyclical attractor includes most of the 16 pure strategies. The most prominent role is for *TFT*, followed in time by *con-C* (and, in lower frequencies, *willing* and *ALLC*), then by a mix of strategies including *ATFT*, *desperate*, *con-D*, *grim*, *ALLD*, and *Pavlov*, to be followed by *TFT* again. After this second peak of *TFT*, a similar mix of

strategies follows as after the first peak, although *ATFT* is absent this time, there is a small peak of *inconsistent*, and there is a longer period of a mix of *ALLD* and *grim* before the cycle starts over again. Except for this last stretch, a small frequency of *MNG* is present throughout. Only a minority of strategies do not play a role in this cyclical attractor: *Acon-D*, *APavlov*, *hopeless*, and *consistent*.

Table S2. ESSs identified by the game theoretical analysis of the IPD and ISD for the alternative payoff configuration. The last column gives the average cooperation level in each equilibrium. Note that there is also a fourth cyclical attractor in the IPD, of which the dynamics are illustrated by Fig. S4. A full description of all strategies can be found in Table 1 of the main text.

<i>Game</i>	<i>ESS</i>	<i>Strategies</i>	<i>Fraction</i>	<i>Cooperation</i>
IPD	1	<i>ALLD</i>	1.000	0.010
	2	<i>grim</i>	1.000	0.013
	3	<i>Pavlov</i>	1.000	0.943
ISD	1	<i>con-D</i>	1.000	0.206
	2	<i>Pavlov</i>	1.000	0.943




Figure S4. A cyclical attractor in the IPD with alternative payoff matrix. Note that some of the colours associated with the strategies are different than in the other graphs.

Figure S5 shows typical behaviour of simulations of the IPD for the altered payoff configuration. These simulations were run for all mechanistic implementations, and also for high (10^{-3}) and low (10^{-4}) mutation rates. It is apparent that in comparison to the original game the periods of stasis are more pronounced, while the periods of fluctuation tend to be much shorter. For example, the 1:1 mapping with per-locus mutation now typically converges to the ESS *Pavlov*, while strong fluctuations were the rule in the original game (Fig. S2). The reduced tendency to cycle is somewhat surprising, since the altered IPD has a cyclic attractor (Fig. S4) while the original game only had two pure-strategy attractors (Table S1).

As in the original model, behavioural architecture and mutation regime strongly affect the outcome and dynamics of evolution. For instance, in the 1:1 mapping, the *Pavlov* equilibrium is a much more common outcome than in the neural network implementation. The *Pavlov* equilibrium was never observed to be invaded by any other strategy in any of the replicate simulations across all implementations. This suggests that given enough time, all simulations should end up in this equilibrium, the waiting time until this happens being determined by the mechanistic implementation. Consistently with this, the *Pavlov* equilibrium was more frequently observed for higher mutation rates. As for the original payoff configuration, the incidence of highly dynamic periods is lower for reduced mutation rates (especially in the neural network implementation). Some of the simulation dynamics bear some resemblance to the cyclical attractor that was identified for this payoff configuration (periods with many subsequent peaks of *TFT*).

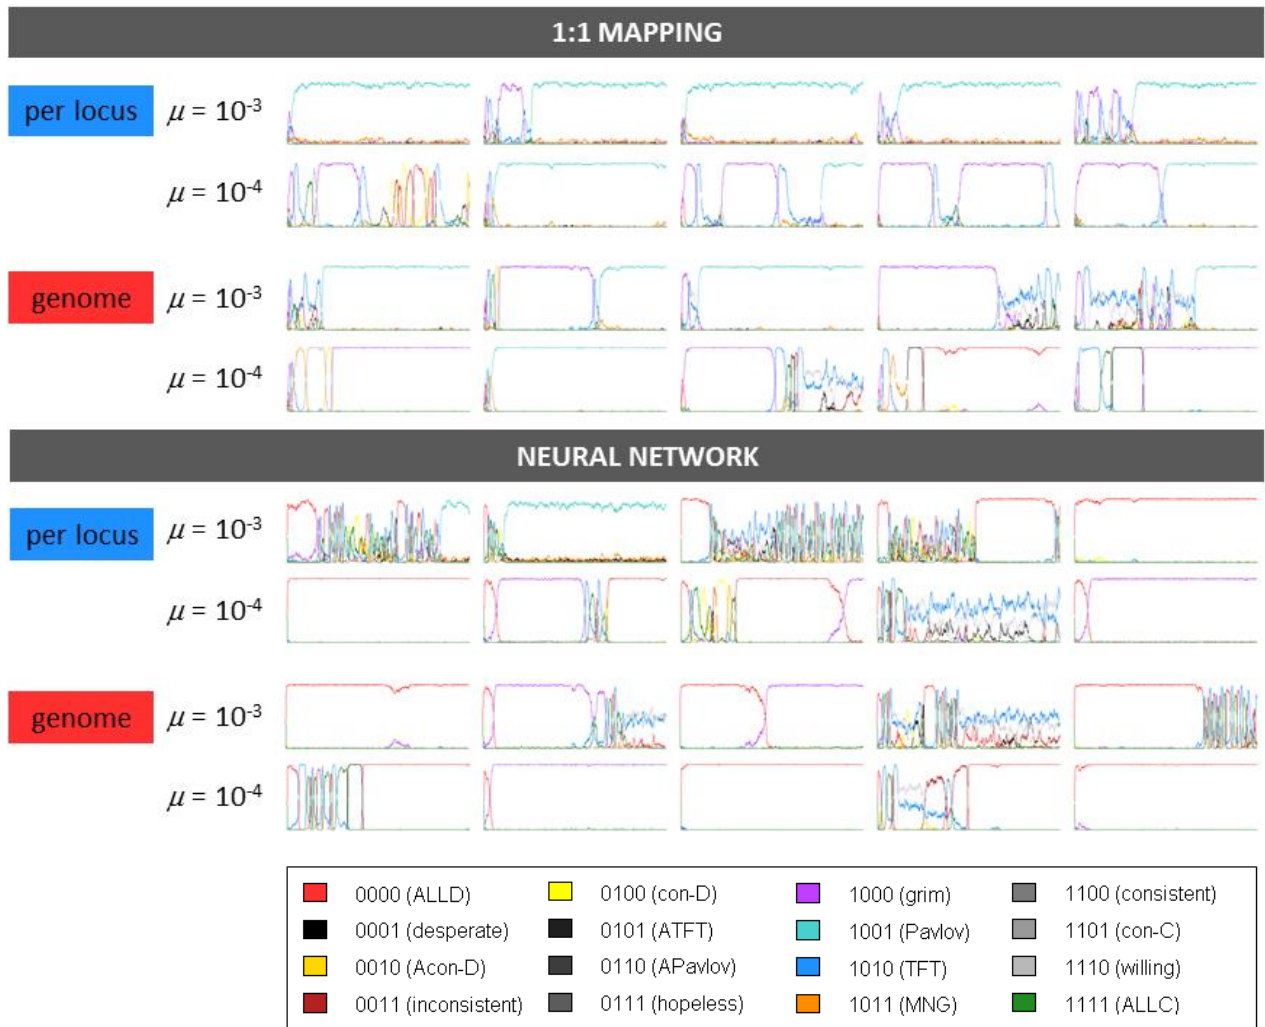


Figure S5. Typical simulation runs for the IPD with altered payoff matrix, for different behavioural architectures, mutation regimes, and mutation rates. As in Fig. S2, representative excerpts of 100 replicate simulation runs are shown for each parameter combination. Graphs of all 100 replicate simulation runs for each parameter combination are available upon request.

Figure S6 shows the equilibria that were attained in the simulations of the ISD with alternative payoff configuration, again for all mechanistic implementations. This altered version of the ISD has two pure-strategy ESSs: *con-D* and *Pavlov* (similar to ESSs 1 and 2 in the original game). From the figure, it is obvious that the mechanistic implementation has a strong effect on the simulation outcome; in the 1:1 mapping, ESS 2 was by far the most common outcome, whereas in the neural network, ESS 1 was more common. The mutation regime also affected the outcome, especially in the neural network

implementation, where ESS 2 was more common for entire-genome mutation than for per-locus mutation. As in the original game, ESS 2 was observed to succeed ESS 1, but never *vice versa*. Accordingly, one would expect to find ESS 2 more frequently in case of a higher mutation rate (since a higher mutation rate should lead to a faster transition from ESS 1 to ESS 2). Indeed, ESS 2 was observed less frequently for the lower mutation rate across all implementations. However, this effect is not very pronounced except for the neural network implementation with entire-genome mutation.




Figure S6. Evolutionary outcome of the simulations of the ISD with alternative payoff configuration, for the different mechanistic implementations and mutation rates. The left bars of each pair show high mutation rates (H, $\mu = 0.001$), bars on the right show low mutation rates (L, $\mu = 0.0001$). Bars indicate the fraction of 100 replicate simulations ending up in each of the three ESSs of the ISD (Table S2) after 100,000 generations.

References

- [a] Van Doorn, GS, Hengeveld GM, Weissing, FJ. 2003 The evolution of social dominance. I. Two-player models. *Behaviour* **140**, 1305-1332.
- [b] Van Damme, E. 1991 *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag.
- [c] Boyd, R, Lorberbaum, JP. 1987 No pure strategy is evolutionarily stable in the repeated Prisoner's Dilemma game. *Nature* **327**, 58-59.
- [d] Maynard Smith J. 1982 *Evolution and the Theory of Games*. Cambridge: Cambridge Univ. Press.
- [e] Selten R. 1983. Evolutionary stability in extensive 2-person games. *Math. Soc. Sci.* **5**, 269-363.
- [f] Hofbauer J, Sigmund K. 1988 *The Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge Univ. Press.
- [g] Weissing FJ. 1991 Evolutionary and dynamic stability in a class of evolutionary normal form games. In: Selten, R. (ed.) *Game Equilibrium Models I. Evolution and Game Dynamics*. Berlin: Springer-Verlag, 29-97.
- [h] Weissing FJ, Van Boven M. 2001 Selection and segregation distortion in a sex-differentiated population. *Theor. Pop. Biol.* **60**, 327-341.