SI APPENDIX

The importance of mechanisms for the evolution of cooperation

Pieter van den Berg & Franz J. Weissing

CONTENTS

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 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.

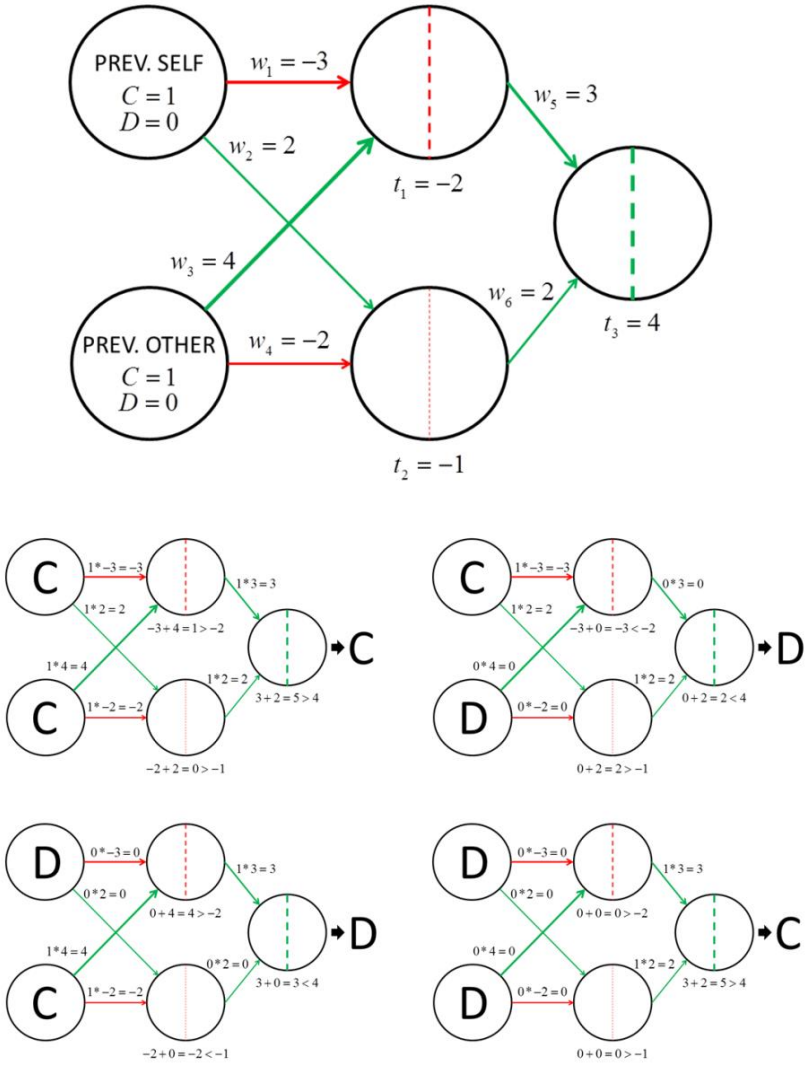

Figure S1. A neural network corresponding to the *Pavlov* strategy. The network on top shows the values of all the weights ($w_1 - w_6$) and thresholds ($t_1 - t_3$); green lines depict positive values, red lines depict negative values, and their thickness indicate their absolute values. The input nodes 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.

Game	*ESS*	*Strategies*	*Fraction*	*Cooperation*
IPD	1	*ALLD*	1.000	0.010
	2	*grim*	1.000	0.013
ISD	1	*con-D*	0.833	0.180
		ALLD	0.167	
	2	*Pavlov*	1.000	0.943
	3	*MNG*	0.968	0.698
		inconsistent	0.022	
		Acon-D	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 $6*10^7$ 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.

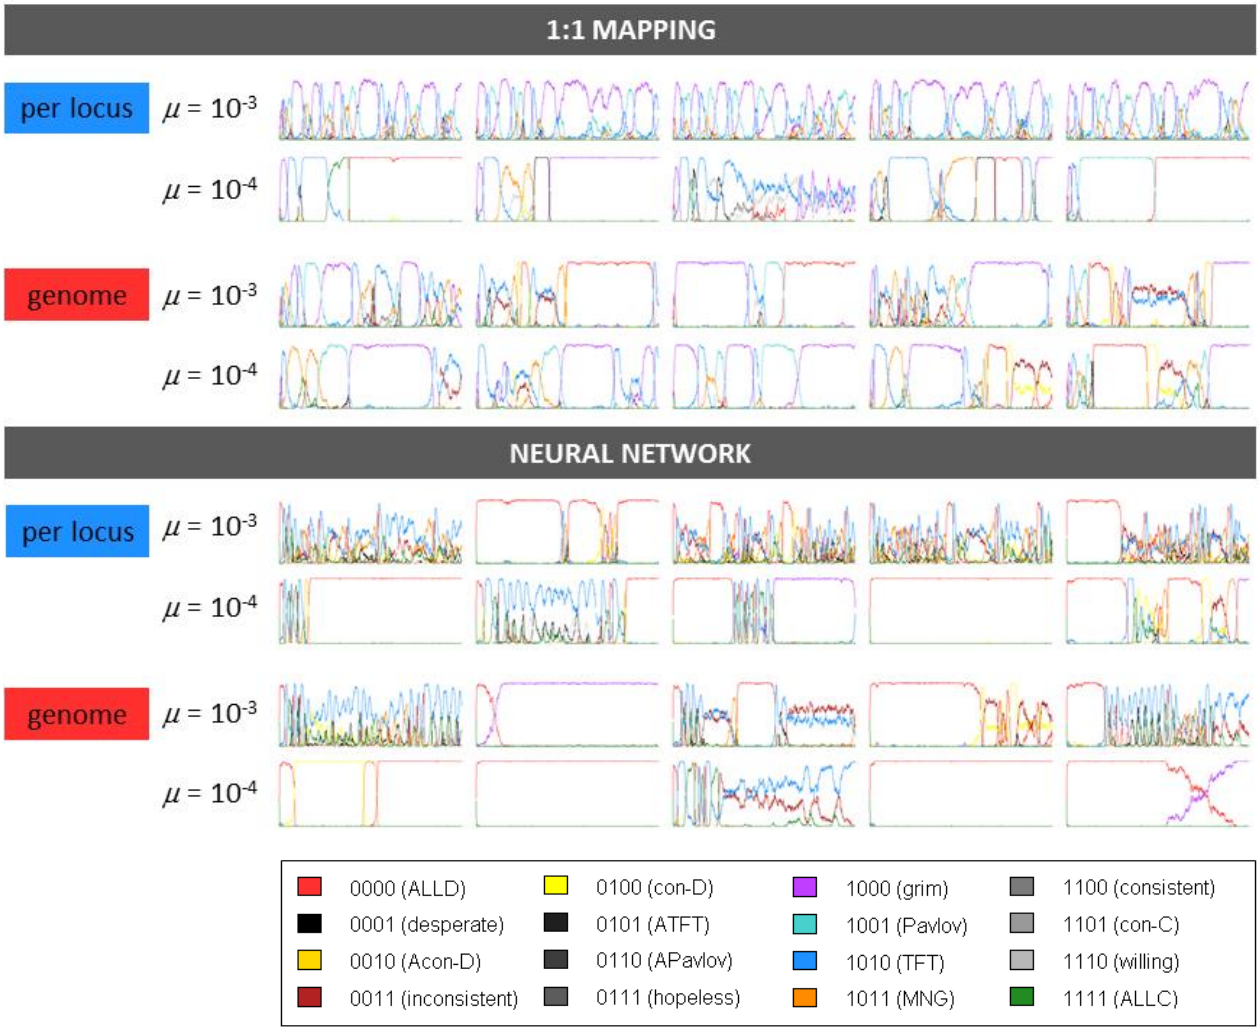

Figure S2. Typical simulation runs for the IPD, for different behavioural architectures, mutation regimes, and mutation rates. For each parameter combination, excerpts of five replicate simulation runs are shown. Each excerpt comprises a time period of 2,000 generations. The excerpts were chosen from 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.

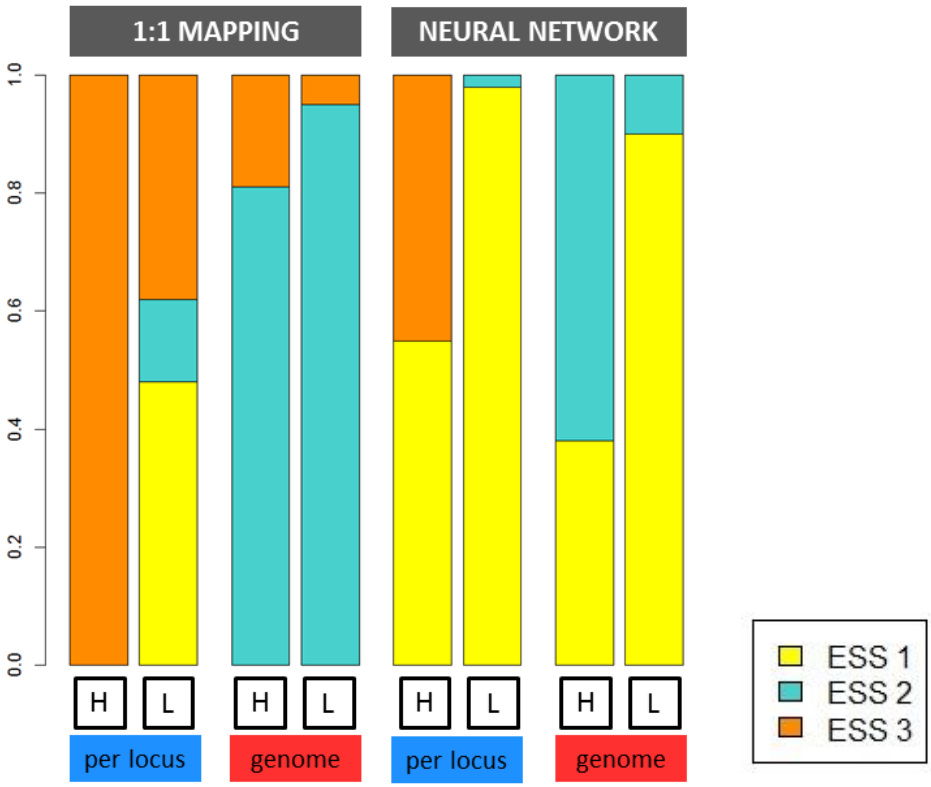

Figure S3. Evolutionary outcome of the simulations of the ISD for the different mechanistic implementations and mutation rates. The left bars of each pair show high mutation rates (H, μ = 0.001), bars on the right show low mutation rates (L, μ = 0.0001). Bars indicate the fraction of 100 replicate simulations 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 he 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.

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

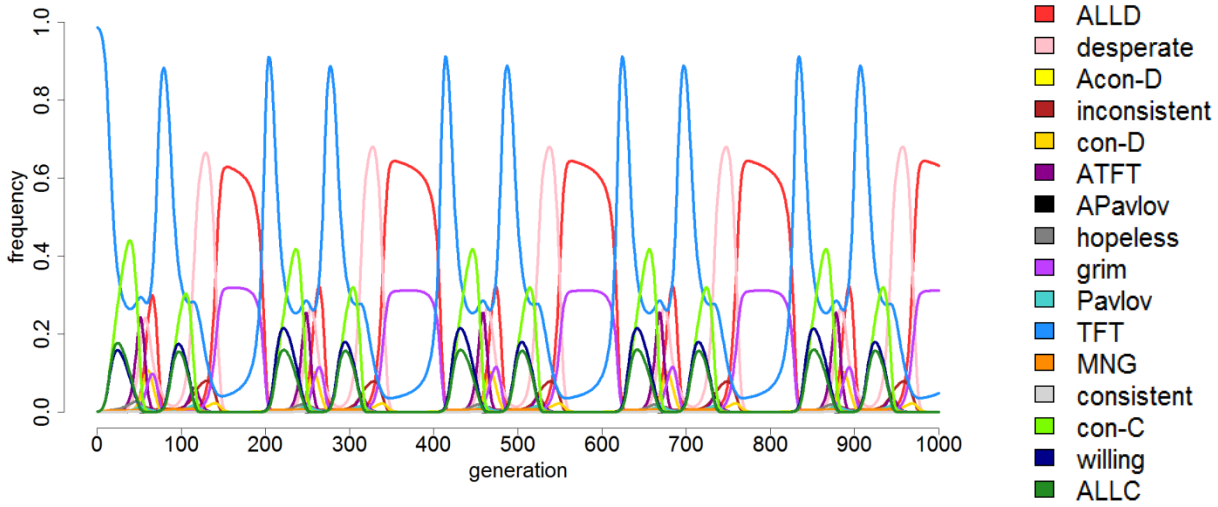

Figure S4. A cyclical attractor in the IPD with alternative payoff matrix. Note that some of the colours associated with the strategies are different than in the other graphs.

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*).

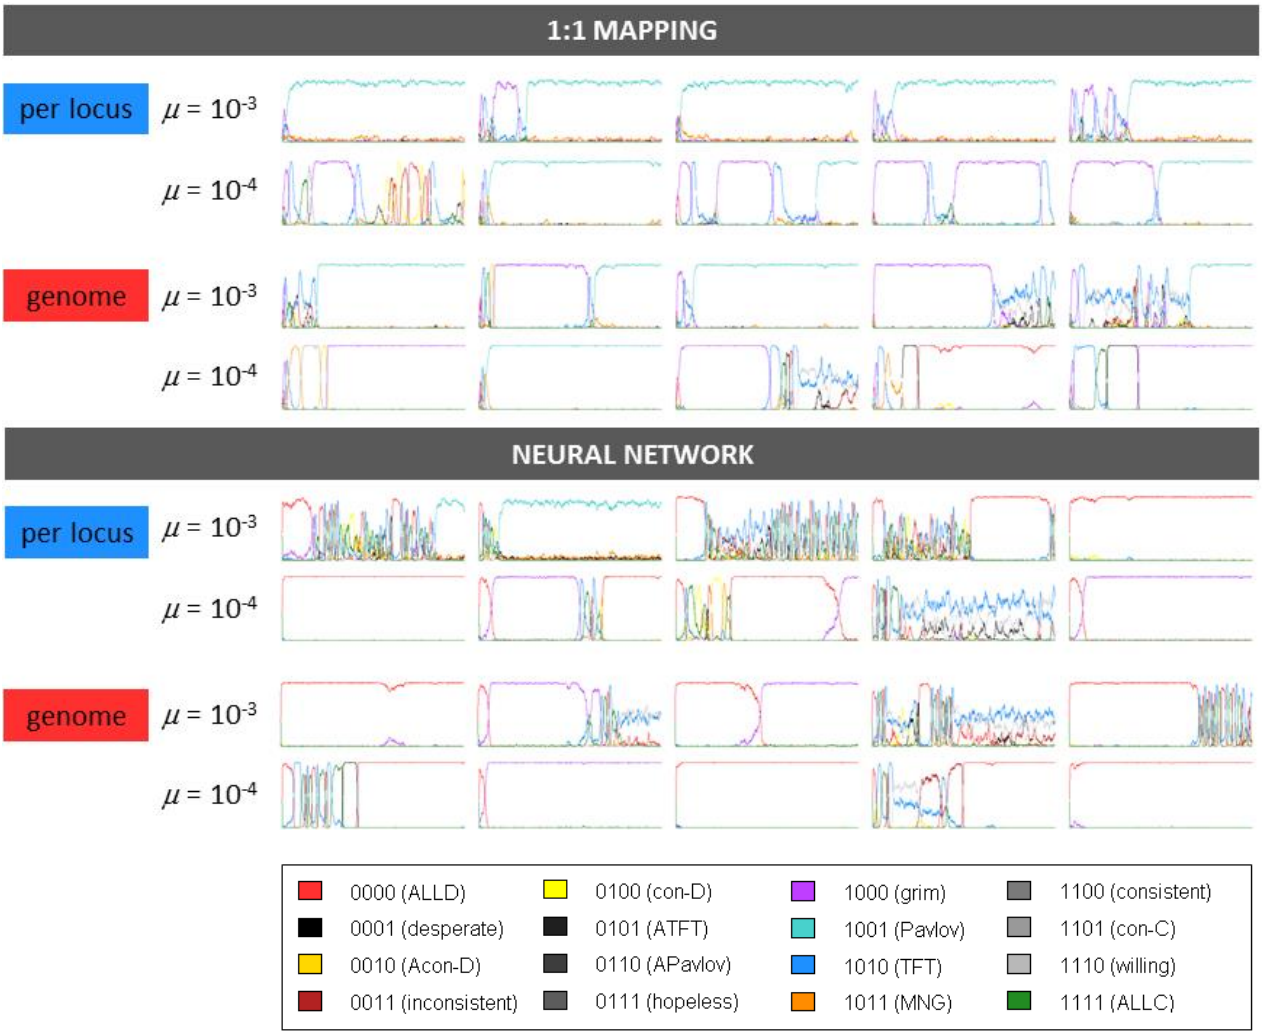

Figure S5. Typical simulation runs for the IPD with altered payoff matrix, for different behavioural architectures, mutation regimes, and mutation rates. As in Fig. S2, representative excerpts of 100 replicate simulation runs are shown for each parameter combination. Graphs of all 100 replicate simulation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

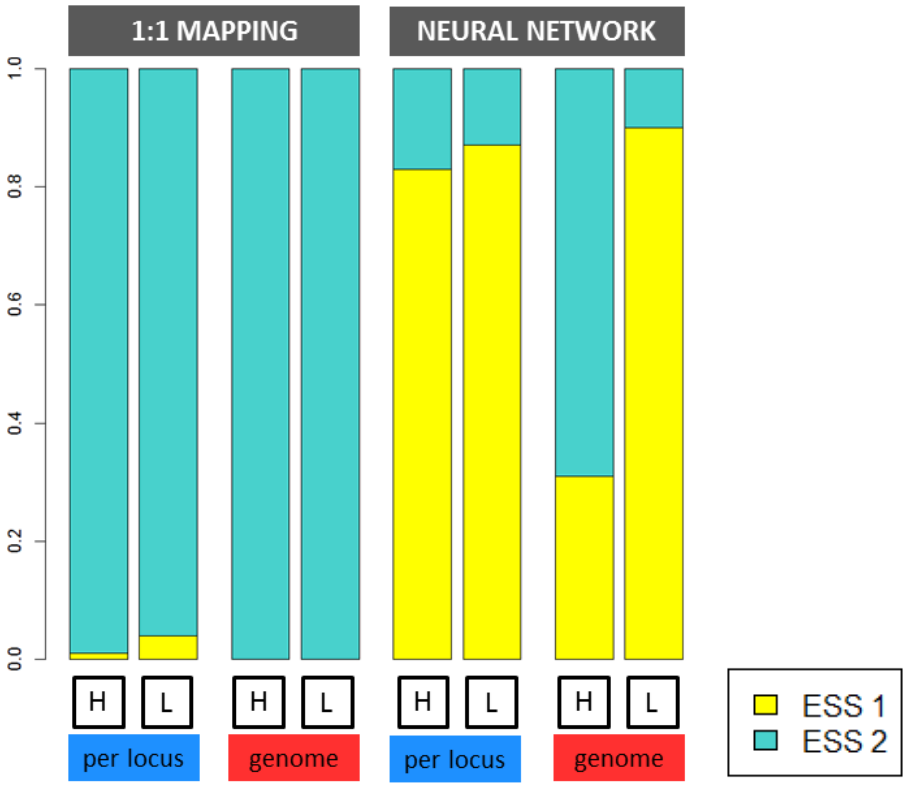

Figure S6. Evolutionary outcome of the simulations of the ISD with alternative payoff configuration, for the different mechanistic implementations and mutation rates. The left bars of each pair show high mutation rates (H, $\mu = 0.001$), bars on the right show low mutation rates (L, $\mu = 0.0001$). 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.