

Cooperation and the Division of Labour

Simon J. Tudge¹, Richard A. Watson¹ and Markus Brede¹

¹ The University of Southampton, Southampton SO17 1BJ, UK
sjt4g11@soton.ac.uk

Abstract

Cooperation is vital for maintaining the integrity of complex life forms. In many cases in nature cooperation manifests itself through constituent parts performing different, but complementary, functions. The vast majority of studies on the evolution of cooperation, however, look only at the special case in which cooperation manifests itself via the constituent parts performing identical tasks. In this paper we investigate a class of games in which the socially optimal behaviour has the property of being heterogeneous. We show that this class of games is equivalent to a region of ST space (the space of normalised two-player games characterised by the ‘sucker’ and ‘temptation’ payoffs) which has previously been dismissed. We analyse, through a simple group selection model, properties that evolving agents would need to have in order to “solve” this dilemma. Specifically we find that positive assortment on pure strategies may lower mean individual payoff, and that assortment on mixed strategies will increase payoff, but not maximise it.

Introduction

Division Of Labour (DOL) is ubiquitous in the biological world. Social insects often have specialised castes for performing individual tasks (Hölldobler and Wilson, 2009). Multicellular organisms exhibit high levels of cell differentiation. Colonial marine invertebrates have differentiated parts which also specialise (Dunn and Wagner, 2006). Even bacteria have been shown to exhibit specialisation (Crespi, 2001). Arguably DOL is one of the major benefits to group living. It has long been recognised that specialisation may result in gains in efficiency; the idea can be traced at least as far back as Adam Smith’s *Wealth of Nations* (Smith, 1776). However, with all group living comes the potential for the emergence of cooperative dilemmas. Whenever a task is broken down into smaller parts the products of the sub-tasks must be shared or distributed. This potentially opens the door to free riders who benefit from the distribution of the products of labour, without contributing to its costs.

There are a growing body of artificial life studies concerning the evolution of the division of labour. Specifically authors have addressed: the mechanisms by which a division of labour can occur (Goldsby et al., 2012), the evolu-

tionary pathway to the emergence of complex internal features (Lenski et al., 2003), the evolution of differentiation in multicellular organisms (Ray and Hart, 1999), the role of gene networks in multicellular development (Joachimczak and Wróbel, 2008) and the evolutionary role of asymmetric cell division (Hotz, 2004).

DOL is also one of the key theoretical ideas behind the major evolutionary transitions research program (Maynard Smith and Szathmari, 1997). A major transition is one in which biological entities which were, preceding the transition, able to replicate as individuals are, after the transition, only able to replicate as part of a larger whole. DOL is likely to be one of the key concepts that leads to a deeper understanding of the major transitions. As increased specialisation develops, individuals become increasingly dependent upon one another, to the point where it is no longer sensible to regard them as functionally independent entities. For example, a potentially defining characteristic of certain types of major transition (i.e. the fraternal transitions (Queller, 1997)) is a reproductive division of labour (Michod, 2006).

Cooperative dilemmas are the class of games in which well-mixed populations of agents evolve to a state which does not maximise mean individual payoff. Theoretical considerations regarding the evolution of cooperation posit a game in which the socially optimal behaviour, for the population, is for every agent to perform the action labelled as cooperate. These set of games are cooperative dilemmas if there exists an ESS which is different from total cooperation. That is that under freely evolving conditions the population is composed either partially or entirely of defectors. Models of the evolution of cooperation then typically consider extensions of the underlying game which result in an increase in the level of cooperation. A common way in which this is achieved is through imposing population structure which leads to positive assortment and hence to an increase in cooperation (see for instance: Nowak and May (1992); Maynard Smith (1964)). Here positive assortment means that like strategies play each other more often than would be expected from random interactions. (For general arguments concerning the role of assortment in the evolution

of cooperation see: Eshel and Cavalli-Sforza (1983); Queller (1985); Dugatkin and Mesterton-Gibbons (1992); Godfrey-Smith (2008)).

Despite the large variety of computational and analytic models along these lines all such studies are built on a common assumption: namely the final optimal state is a homogeneous one in which all individuals play the same strategy. However, many situations in nature can be said to be in cooperative states, but between individuals or components which are not exhibiting homogeneous behaviours. To the best of the authors' knowledge none of these evolutionary game-theoretic investigations have considered situations in which a heterogeneous final state is desirable.

This paper firstly identifies the class of games in which a mixed state is socially optimal, i.e. games in which a division of labour may evolve. We show that these games are related to the conventional cooperative dilemmas. We then go on to present two models to illustrate some key points. The first model challenges the assumption that positive assortment on pure strategies will lead to an increase in the population's mean payoff. The second model extends this by introducing the additional assumption of mixed strategies. The model shows that positive assortment on mixed strategies does lead to an increase in the population's mean payoff. Finally we sketch some further theoretical considerations which show that, although positive assortment on mixed strategies does lead to an increase in payoff, it is not the highest payoff that can be reached under any circumstance. Specifically in order to maximise average payoff individuals would have to control not just the frequency of strategies, but the frequency of interactions within the population. In this case a negative assortment on social strategies is optimal for the population; however, we show that it is not evolutionarily stable. In order for this optimal configuration to be stable it is necessary to have a higher level positive assortment on genotypes which provides a lower level negative assortment of phenotype/social strategy.

Division of Labour Games

We now outline a formalism which enables us to think about the division of labour in the simplest non-trivial case.

Consider a situation in which individuals meet and perform one of two tasks: A or B. Each task bestows a benefit to both of the individuals involved in the interaction. The benefits are given by b_A and b_B respectively. Each individual must bear the cost of their performed task themselves. Costs are given by c_A and c_B . However, if both individuals perform the same task the cost of that task is shared between them. In addition, there is a synergistic benefit which is the benefit of having both tasks performed together: δ . We consider the cases in which A has a higher cost but also a higher benefit than B, i.e. $c_A > c_B$ and $b_A > b_B$.

An example of the situation described above might go as follows. Two human individuals living in the same tribe may

perform one of two tasks. Task A is to go and hunt for meat. Task B is to build a fire. Hunting comes at an extra cost to the individual either because it requires more energy or because it is inherently riskier. The benefit of hunting is meat. Building the fire has a lower, but non-zero, cost. The benefit of building the fire is warmth. We assume that the meat is more valuable than the warmth, but that both tasks provide some benefit in isolation. In this instance the synergistic benefit, δ , is that of having cooked meat. It is the benefit above and beyond that of the sum of the two benefits in isolation. We are assuming here that the benefits are non-excludable, that is that the hunter could not stop the fire builder from taking meat, and vice versa. This cartoon is an aid to understanding; the essential features of the situation are represented via the payoff matrix:

	A	B
A	$b_A - c_A/2$	$b_A + b_B + \delta - c_A$
B	$b_A + b_B + \delta - c_B$	$b_B - c_B/2$

Given that games have two arbitrary degrees of freedom we will assume that: $b_A - c_A/2 = 1$ and $b_B - c_B/2 = 0$. We can thus rewrite the above payoff matrix as:

	A	B
A	1	$1 - r + \delta$
B	$1 + r + \delta$	0

Where $r = \frac{1}{2}(c_A - c_B)$. This reduces to the 1 dimensional parameterisation of the snowdrift game if $\delta = 0$ (Hauert, 2004). Conceptually we may also think of the story behind the snowdrift game as the special case in which task B is the task of doing nothing with no benefit and no cost.

We can trivially see that this game represents a region of ST space using $S = 1 - r + \delta$ and $T = 1 + r + \delta$. For an explanation of ST space see Santos et al. (2006).

Note that r is the difference in cost between performing the two tasks, and can thus be thought of as parameterising the severity of the dilemma. δ represents the synergistic benefits of having both tasks performed. $\delta > 0$ corresponds to the region $S + T > 2$.

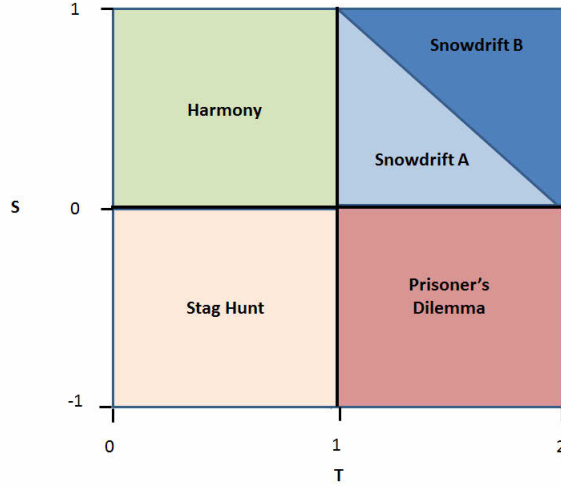


Figure 1: The location of all games in ST space. Previously the top right quadrant was simply referred to as the snowdrift game. Here we split this region into two. Snowdrift A corresponds to games which are snowdrift games but not DOL games. Snowdrift B refers to snowdrift games which are also division of labour games. The remainder of this paper is concerned with the region labelled snowdrift B.

Cooperative Dilemmas

The previous section introduced a class of games which are formally equivalent to the region of ST space in which $S + T > 2$. We shall refer to these games as Division of Labour games. Snowdrift games are defined by $0 < S < 1$ and $T > 1$. We shall focus our investigation on the class of games which are both snowdrift games and DOL games (note that neither one implies the other).

ST space was conceived of in order to systematically investigate all classes of cooperative dilemmas. Some authors (see for instance Macy and Flache (2002)) specifically exclude the region $S + T > 2$ from the definition of cooperative dilemmas. We find this exclusion somewhat artificial. The essence of a cooperative dilemma is a situation in which evolution leads to a state which does not maximise mean individual payoff. It just so happens that in conventional cooperative dilemmas social welfare is maximised by every agent cooperating, but this is by no means an essential part of the argument. Let us define the Socially Optimal Frequency (SOF) as the frequency of cooperate (or type A in the language of DOL games) which maximises the mean payoff of the population under well mixed conditions. We then define a cooperative dilemma as one in which $\text{SOF} \neq \text{ESS}$.

We now need to derive an equation for the SOF in terms of S and T . To do this note that the mean fitness of the population is given by:

$$\bar{f} = f_c p_c + f_d p_d \quad (1)$$

We consider a fixed population size so that $p_d = 1 - p_c$. Fitnesses are given by $f_c = p_c + S p_d$ and $f_d = T p_c$. We then arrive at a formula for average fitness in terms of density of cooperators (ρ):

$$\bar{f} = \rho(S + T + \rho(1 - S - T)) \quad (2)$$

The SOF is the maximum of this function for $\rho \in [0, 1]$. It is straight forward to prove that:

$$\text{SOF} = \begin{cases} \frac{S+T}{2(S+T-1)} & S+T > 2 \\ 1 & S+T \leq 2 \end{cases} \quad (3)$$

The ESS of the snowdrift game is (Nowak, 2006a):

$$\text{ESS} = \frac{S}{S+T-1} \quad (4)$$

for DOL games which are also snowdrift games the ESS is only equal to the SOF in the very special instance in which $S = T$. Thus, by our slightly broader definition, DOL games are cooperative dilemmas.




Figure 2: Left: The equilibrium frequency of cooperate. Right: the SOF, note that in the top right hand corner this is not equal to 100% cooperation.

Model

In this section we demonstrate that positive assortment on pure strategies is only sufficient in allowing populations to reach the SOF in the non-generic case in which $\text{SOF} = 1$. We go on to show that positive assortment is still one of the key elements in allowing populations to reach the SOF. However, in this case the assortment must be on something other than pure social strategies.

We implement two generational GAs (labelled model I and II) to illustrate some key points. In both cases we construct a scenario in which one can exogenously control the level of assortment in a population of evolving individuals and measure the total payoff in the population. Model I serves as a control for model II. With model I we allow only pure strategies which are only able to perform one of two tasks, A or B, for the entirety of their lifetime. In model II we lift this assumption and allow for mixed strategies. Specifically a genotype specifies not a task A or B, but a probability, $p \in [0, 1]$, which determines how often task A is

performed. Apart from this difference models I and II have the same underlying structure.

The models consist of two distinct phases, a group phase and a population phase. Rounds of the game during one generation occur in τ discrete time steps; the first X time steps within groups, the remaining $\tau - X$ time steps within the fully mixed population. If f_G is fitness acquired within the group phase and f_P fitness acquired within the population phase then total fitness, f_T , is given by:

$$f_T = \frac{X}{\tau} f_G + \frac{\tau - X}{\tau} f_P \quad (5)$$

There are N_G groups consisting of g players in each group. Individuals acquire fitness over both stages of the generation. N_G individuals are chosen at the end of the population stage via fitness proportionate selection and go on to form new groups. The founding individuals immediately replicate $g - 1$ times so that the groups are composed of g clonal individuals. Figure 3 shows schematics for the two models.




Figure 3: Schematics for models I and II respectively.

We consider the model in two different manners. First of all, we model the situation through numerical integration of the relevant replicator equation (see appendix A). This corresponds to infinite populations without mutations. Secondly, we model the system via an agent based simulation with finite population and mutations. The two approaches show good agreement in final results.

In the agent-based model there is mutation. Mutation

leads to one of the individuals in the group stage being genetically different to their parent.

In model I we model mutation by allowing an A to create a B, and vice versa, with a probability $\mu = 1 \times 10^{-2}$. In model II, with probability $\mu = 5 \times 10^{-2}$, an individual is born with a value of p which differs from its parent by an amount chosen from the random uniform distribution $[-0.1, 0.1]$. Mutations are capped to physically meaningful ranges (i.e. between 0 and 1) if they mutate outside of this range.




Figure 4: Top: The mean value of p at equilibrium for a range of values of α in model II. The lower dotted line is the ESS and the upper one the SOF. Bottom: Fitness at ESS for the two models. The solid line is generated from the predictions of the analytic model, points from the agent based model. Clearly agents in model II are at an advantage over those in model I. $(r, \delta) = (0.5, 0.5)$.

This group structuring model allows us to fine-tune the level of assortment on the population. It bears some similarities to the hay-stack model (Maynard Smith, 1964). However, within a group there is no selection, as all members are clonal. The groups are formed from a founder and serve only to limit the interactions of individuals to a certain, non-random, subset of the population. In this sense the models bear some conceptual similarities to the ones discussed in Godfrey-Smith (2008). The qualitative results would be repeatable with any of the standard repertoire of “evolution of

cooperation” models (Nowak, 2006b). This particular model is not chosen for biological realism, but because the essential property of population structure leading to positive assortment is completely transparent. The models have the convenient feature of being able to tune the level of population structure via the parameter $\alpha = x/\tau$.

Previously the distinction between pure and mixed strategies has not been important. In a snowdrift game the frequency of cooperate at the ESS can represent either the frequency of pure strategy cooperators within the population or the average value of p (the probability of cooperating) within a population of mixed strategy individuals (Maynard Smith, 1982, p.17). These two models show that this distinction is important when playing a game in which $\text{SOF} \neq 1$.

In model I there is a one-to-one correspondence between the genotype and the social strategies A and B. Group structure provides only positive assortment on pure strategies. Assortment is needed for the evolution of cooperation. However, positive assortment leads to groups comprised only of one type.

It is important to realise that model II provides assortment on mixed strategies rather than on the pure phenotypes A and B. This is crucial to the following results.




Figure 5: The mean fitness at ESS. Left column is for pure strategies, as in model I, right column is for mixed strategies, as in model II. Going down the page the figures correspond to increasing levels of α . Top: $\alpha = 0$, middle: $\alpha = 0.25$ and bottom: $\alpha = 0.5$.

Figure 4 shows the results of the two versions of the model. There are 125 groups composed of 5 individuals each. The simulation is run for 5×10^6 generations and the

average value of p (probability of playing strategy A) over the entire population is recorded. We find that groups in model II evolve towards the SOF for larger values of α (the degree of population structure). As expected with $\alpha = 0$ the population simply evolves to the ESS in both models.

We plot the relative fitness for models I and II for increasing α to illustrate the fact that mixed strategies are at an evolutionary advantage over pure ones. Note that by construction pure strategies only interact with types of the same pure strategy within groups (in model I). In the case where the socially optimal solution was pure cooperate the ability to form pure groups is sufficient to solve the dilemma. However, in general these are a special type of game. For DOL games, in which a mix of strategies is desirable, mixed strategies can outperform pure ones.

Figure 5 shows the fitness at ESS for models I and II for all games parameterised via r and δ . Interestingly positive assortment on pure strategies can actually be detrimental to the population’s payoff if δ (the synergistic benefit to heterogeneous behaviours) is sufficiently high. The higher the synergistic benefit is the greater the advantage of having mixed strategies.

Group Phenotypes

Groups composed of mixed strategies do not maximise social welfare. The reason for this is that they are unable to control their internal structure or organisation. In this section we formalise this point. We leave a detailed specification and analysis of a model for a forthcoming work.

In a well-mixed population the parameter ρ (frequency of type A) characterises the state of the system. However, if interactions are not random then ρ on its own is insufficient. We also need to know the frequency of the different types of within-group interactions. In principle there are three types of interactions: (A-A, A-B and B-B). If the total number of interaction of all types is fixed, then knowing the fractional density of each type will specify the state of the group. We will denote these three densities as ϕ_{AA} , ϕ_{AB} and ϕ_{BB} . However, it is sufficient to know only one of these. Let us then use ϕ_{AB} and drop the subscript. We shall define the *group phenotype* as a point in the space (ρ, ϕ) . The following formulae show how the densities of all types of interaction can be found from these two variables.

$$\phi_{AA} = \rho - \frac{1}{2}\phi \quad (6)$$

$$\phi_{AB} = \phi \quad (7)$$

$$\phi_{BB} = 1 - \rho - \frac{1}{2}\phi \quad (8)$$

Notice that ϕ is confined within certain ranges based on ρ . Specifically $0 \leq \phi \leq 2\text{Min}\{\rho, 1 - \rho\}$. ϕ is equivalent to certain measures of linkage disequilibrium (see for instance Hartl and Clark (1998)). For interesting parallels between population genetics and the evolution of cooperation

see Gardner et al. (2007).

The previous section considered the cases in which groups only had control over the parameter ρ and had no way of controlling the internal composition of within group interactions. In this case they were confined to have a value of

$$\varphi = \varphi^{(R)} = 2\rho(1 - \rho) \quad (9)$$

where R stands for random. Thus the SOF (Socially Optimal Frequency) corresponds to the fittest group with random internal interactions. Let us also define the Optimal Group Phenotype (OGP) to be the point in (ρ, φ) space which maximises average fitness. To see this, note that average fitness is given by:

$$\bar{f} = \rho + \frac{1}{2}(S + T - 1)\varphi \quad (10)$$

for games in (S, T) space. The OGP is then given by:

$$OGP = \begin{cases} (1, 0) & S + T < 2 \\ (1/2, 1) & S + T > 2 \end{cases} \quad (11)$$

thus in DOL games a group which maximises the amount of A-B links at the expense of all other types of interactions is the one which maximises group fitness.




Figure 6: Left: The average fitness in terms of ρ and φ , for a certain game in which $R + T > 2$. Right: the absolute difference between the fitness of the two phenotypes A and B (assuming that the group is composed of pure strategy individuals). The dotted line marks $\varphi^{(R)}$. Circled points correspond to A: The ESS under well mixed conditions, B: The SOF, C: The optimal point for pure groups and D: the OGP. Note: if $S + T < 2$ B, C and D coincide. The fact that stability is non-zero at points B and D shows that these situations are fundamentally unstable without the addition of extra assumptions.

With reference to figure 6 model I of the previous section was only able to evolve to point C, model II was able to do

better and reach point B. However, neither model produced groups who were able to reach point D. This is the subject of a forthcoming work in which we investigate the effects of developmental or aggregational processes.

Discussion

Kant said:

“Act only according to that maxim whereby you can, at the same time, will that it should become a universal law.”

In a fully assorted population Kant’s principle is not only morally commendable, but it is also entirely sensible. Given that you will only meet individuals who are the same as you it makes sense to perform social actions which are beneficial to be on the receiving end of. Thus it would seem that positive assortment is *the* answer to the evolution of cooperation.

On the other hand gains from specialisation occur via a collection of *different* types of individuals. We have seen that a division of labour game may be a cooperative dilemma. There are two needs which seem to be fundamentally at odds with each other: firstly, the need for positive assortment to alleviate the cooperative dilemma, and secondly, the need for negative assortment in order to gain from specialisation. This is the fundamental problem of the evolution of the division of labour. How does nature have her cake and eat it? That is how does evolution create the positive assortment necessary to alleviate the cooperative dilemma, but at the same time maintain the diversity needed to benefit from a division of labour?

It could be argued that many interesting and complex aspects of the biological world are about solving this problem. Phenotypic plasticity is a way in which a social strategy is able to become decoupled from the genotype which underlies it (Gavrilets, 2010). Thus we can have assortment on genotype without assortment on phenotype (as in model II), which goes some way to alleviating the problem of the division of labour. This is the key point which the models presented here attempt to illustrate. One way of expressing this would be to say that the social strategy has become de-Darwinised (*sensu* Godfrey-Smith (2009)). In division of labour games the optimal configuration involves As interacting with Bs to the exclusion of all other types of interaction. However, in this situation the fitness acquired by the phenotype B will always outweigh that acquired by A (because in our framework $T > S$). Thus the optimisation of the higher level entity, the group, is in direct conflict with that of the lower level entities, the individuals. The only way in which higher level optimisation can occur is if selection does not act directly on the frequency of the constituent types A and B. This is what we mean by de-Darwinisation. A potential way for this to occur is through a genotype-phenotype map which is not one-to-one (i.e. a genotype may specify more

than one phenotype). In this case, although social interactions lead to Bs having a higher fitness, the configuration is sustainable because natural selection does not “see” the phenotypes A and B, it only “sees” genotypes which specify certain frequencies and organisations of social strategies. We see the concept of de-Darwinisation as a powerful conceptual tool for understanding the emergence of higher levels of biological organisation.

An ALife approach will doubtless be one of the key theoretical tools in our quest to understand biological organisation. Simulation is necessary not only because the processes of interest are obscured by time, but also because we only have one truly independent example of life. What we really want to know is which aspects of biology are contingent on the particulars of our bio-chemistry, and which are profound consequences of the logic of natural selection. This paper has attempted to add to the small, but growing number, of ALife studies which tackle the question of the division of labour and internal organisation. We have laid down ground work for a systematic investigation of the ultimate causes of the evolution of internal differentiation and organisation.

Acknowledgments

This work was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1). Thanks to Paul Ryan and Patrick Doncaster for their useful comments. Also to Charlott Rittmeister for her proof reading.

Appendix A

We sketch the solution to the model via the replicator equation formalism, for the case of mixed strategies playing within variable levels of group structure.

Individuals have strategies which are specified via a probability p . They play strategy A with probability p , and therefore B with probability $1 - p$. An individual with strategy p who interacts with another individual with strategy q receives an expected payoff of:

$$F[p, q] = pqR + p(1 - q)S + (1 - p)qT + (1 - p)(1 - q)P \quad (12)$$

Selection acts on p and thus the population is specified by the function $\rho(p)$ which is a 1D function which specifies the density of the population playing a strategy for every value of $p \in [0, 1]$. The systems dynamics are specified by the replicator equation:

$$\dot{\rho}(p) = \rho(p) (f(p) - \bar{f}(p)) \quad (13)$$

where $f(p)$ is the fitness of the individuals for a given p , and $\bar{f}(p)$ is the mean fitness of the population.

The fitness of any strategy comprises of two parts. The fitness gained in the group phase, and the fitness gained in the population phase. Call these fitness f_G and f_P respectively. In the absence of mutations strategies always play with like

strategies in the group phase, thus:

$$f_G = F[p, p] \quad (14)$$

in the population phase strategies play with every other strategy. The average payoff is given by the strategy they would have received from playing a hypothetical average individual. That is:

$$f_P = F[p, \bar{p}] \quad (15)$$

where \bar{p} is the average value of p in the population.

Total fitness is thus:

$$f(p) = T_G F[p, p] + T_P F[p, \bar{p}] \quad (16)$$

we normalise by saying that $T_G + T_P = \tau$, i.e. that the whole cycle happens over τ units of time. Let $\alpha = T_G/\tau$. By dividing by τ we arrive at:

$$f(p) = \alpha F[p, p] + (1 - \alpha) F[p, \bar{p}] \quad (17)$$

We thus have a fitness defined for every possible strategy, which can be used to model the situation by means of the replicator equation.

Population dynamics follow from numerical integration of equation 13.

References

- Crespi, B. J. (2001). The evolution of social behavior in microorganisms. *Trends in ecology & evolution*, 16(4):178–183.
- Dugatkin, L. A. and Mesterton-Gibbons, M. (1992). Cooperation Among Unrelated Individuals: Evolutionary Factors. *The Quarterly Review of Biology*, 67(3):267–281.
- Dunn, C. W. and Wagner, G. P. (2006). The evolution of colony-level development in the Siphonophora (Cnidaria:Hydrozoa). *Development genes and evolution*, 216(12):743–54.
- Eshel, I. and Cavalli-Sforza, L. L. (1983). Assortment of encounters and evolution of cooperativeness. *Proceedings of the National Academy of Sciences*, 79(4):1331–1335.
- Gardner, A., West, S. A., and Barton, N. H. (2007). The relation between multilocus population genetics and social evolution theory. *The American naturalist*, 169(2):207–26.
- Gavrilets, S. (2010). Rapid transition towards the Division of Labor via evolution of developmental plasticity. *PLoS computational biology*, 6(6).
- Godfrey-Smith, P. (2008). Varieties of Population Structure and the Levels of Selection. *The British Journal for the Philosophy of Science*, 59(1):25–50.

- Godfrey-Smith, P. (2009). *Darwinian Populations and Natural Selection*. OUP Oxford.
- Goldsby, H. J., Dornhaus, A., Kerr, B., and Ofria, C. (2012). Task-switching costs promote the evolution of division of labor and shifts in individuality. *Proceedings of the National Academy of Sciences of the United States of America*.
- Hartl, D. L. and Clark, A. G. (1998). *Principles of Population Genetics*. Sinauer Associates.
- Hauert, C. (2004). Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature*.
- Hölldobler, B. and Wilson, E. (2009). *The Super-Organism*. Norton.
- Hotz, P. (2004). Asymmetric cell division in artificial evolution. *Evolutionary Computation, 2004. CEC2004. . . .*
- Joachimczak, M. and Wróbel, B. (2008). Evo-devo in silico: a model of a gene network regulating multicellular development in 3D space with artificial physics. *Artificial Life XI: Proceedings of the . . .*, pages 297–304.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–44.
- Macy, M. W. and Flache, A. (2002). Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences of the United States of America*, 99 Suppl 3:7229–36.
- Maynard Smith, J. (1964). Group selection and kin selection. *Nature*.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.
- Maynard Smith, J. and Szathmary, E. (1997). *The Major Transitions in Evolution*. OUP Oxford.
- Michod, R. E. (2006). The group covariance effect and fitness trade-offs during evolutionary transitions in individuality. *Proceedings of the National Academy of Sciences of the United States of America*, 103(24):9113–7.
- Nowak, M. A. (2006a). *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press.
- Nowak, M. A. (2006b). Five rules for the evolution of cooperation. *Science (New York, N.Y.)*, 314(5805):1560–3.
- Nowak, M. A. and May, R. (1992). Evolutionary games and spatial chaos. *Nature*.
- Queller, D. C. (1985). Kinship, reciprocity and synergism in the evolution of social behaviour. *Nature*, 318.
- Queller, D. C. (1997). Cooperators since life began. *Quarterly Review of Biology*.
- Ray, T. and Hart, J. (1999). Evolution of differentiated multi-threaded digital organisms. *Proceedings 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems.*, 1:1–10.
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. (2006). Cooperation prevails when individuals adjust their social ties. *PLoS computational biology*, 2(10):e140.
- Smith, A. (1776). *Wealth of Nations*. Hayes Barton Press.