

Kin Selection, Group Selection, and the Varieties of Population Structure

Jonathan Birch

Abstract: Various results show the ‘formal equivalence’ of kin and group selectionist methodologies, but this does not preclude there being a real and useful distinction between kin and group selection processes. I distinguish individual- and population-centred approaches to drawing such a distinction, and I proceed to develop the latter. On the account I advance, the differences between kin and group selection are differences of degree in the structural properties of populations. A spatial metaphor (‘ K - G space’) provides a useful framework for thinking about these differences: kin and group selection may be conceptualized as large, overlapping regions of K - G space. I then consider some implications of the account, defend it from possible objections, and further argue that the structural features characteristic of both kin and group selection may recur at multiple levels of biological organization.

1	Introduction	2
2	Equivalence Results and Their Limitations	3
2.1	An example of an equivalence result	3
2.2	Limitations	6
3	Individual- and Population-Centred Approaches	8
4	Two Influences: Hamilton and Godfrey-Smith	11
5	K and G	14
5.1	K	15
5.2	G	16
5.3	K - G space	18
6	The $rb \neq 0$ Requirement	20

7	Levels of Organization	23
8	The Key Substantive Questions	24

1 Introduction

The relationship between kin selection and group (or multi-level) selection is a longstanding source of controversy in the social evolution literature. In earlier debates, biologists tended to regard kin and group selection as rival empirical hypotheses (Maynard Smith, 1964, 1976; Dawkins, 1982). But many biologists now regard them as ‘formally equivalent’ approaches, and see this equivalence as implying that they are not competing empirical hypotheses after all (Marshall, 2011).¹ Although there are high-profile dissenters from this equivalence claim², including Martin Nowak and Edward O. Wilson, it seems to be endorsed by a majority of social evolution theorists.³

Yet the debate has long been hampered by insufficient attention to the distinction between statistics and causality (Birch and Okasha, 2015; Okasha, 2016). It is crucial to distinguish between the formal equivalence of two statistical descriptions of change and the causal equivalence (or otherwise) of two types of selection process responsible for change. The former does not imply the latter. Indeed, my claim in this article is that, although there is an important sense in which kin and group selection are formally equivalent when conceived as statistical descriptions of change, there is a real and useful—but not sharp—distinction between kin and group selection conceived as causal processes responsible for change. The key differences lie in their commitments regarding population structure.

Here is the article in outline. In the next section, I consider the ‘formal equivalence’ results mentioned above, explaining why these results are compatible with there being a biologically meaningful distinction between kin and group selection. In Section 3, I set out two ways of making sense of this distinction. One ap-

¹Here, and throughout the article, the type of group selection I have in mind is ‘MLS1’ in the terminology of Heisler and Damuth (1987). In other words, it is group selection in which the fitness of a group is defined as the average fitness of its members, rather than the number of offspring groups it produces. I use the term ‘group selection’ in preference to ‘multi-level selection’ because I see both kin and group selection as processes that can occur at multiple levels of organization (see Section 7).

²For example, Traulsen (2010); Nowak *et al.* (2010) and van Veelen *et al.* (2012) depart from this view in various ways. Sober and Wilson (1998) depart from it by regarding kin selection as a special case of group selection. As will become clear in due course, I do not regard either process as a special case of the other.

³For statements of the consensus view, see Wenseleers *et al.* (2010); Gardner *et al.* (2011); Marshall (2015).

proach, developed by Samir Okasha (2016), locates the difference in the causal path at the individual level between an organism’s genotype and its fitness. I highlight some problems with Okasha’s approach that, although not fatal, motivate the development of an alternative. My favoured alternative, which I call the ‘population-centred’ approach, locates the difference in the structural features of populations.

The rest of the article pursues the population-centred approach. In Section 4, I draw inspiration from two sources: W. D. Hamilton’s (1975) views on the relationship between kin and group selection, and Peter Godfrey-Smith’s (2006; 2008) recent work on the varieties of population structure. Section 5 combines these influences into a positive proposal. The intuitive idea is that kin selection occurs in populations that are structured such that relatives tend to interact differentially, whereas group selection occurs in populations in which there are stable, sharply bounded and well-integrated social groups at the relevant grain of analysis. Some populations have both features, but it is possible for one to occur without the other. Since these structural features are matters of degree, a spatial metaphor (‘*K-G* space’) is useful for thinking about the distinction.

The account also requires that $rb \neq 0$ as a precondition for both kin and group selection; in Section 6, I explain and defend this requirement. In Section 7, I discuss the relationship between the kin/group selection distinction and levels of biological organization, arguing that both types of selection process can occur at multiple levels. In Section 8, I conclude by setting out the key questions at stake, by the lights of my account, when we ask whether a process is one of kin selection or group selection.

2 Equivalence Results and Their Limitations

2.1 An example of an equivalence result

The best-known argument for the ‘formal equivalence’ of kin and group selection involves comparing a generalized version of Hamilton’s rule (Queller, 1992a) with the multi-level version of the Price equation (Price, 1972), and noting that both provide correct conditions for positive gene frequency change in a group-structured population, given some fairly minimal assumptions. I will not recount the details of the argument here (for details, see Marshall 2011; Frank 2013; Birch and Okasha 2015; Okasha 2016), but I will briefly explain the source and nature of the equivalence result.

The route to the generalized version of Hamilton’s rule, in the case of a social trait controlled by a single allele, involves mathematically decomposing the overall change in allele frequency between ancestral and descendant populations (such as

two consecutive generations of a discrete generations model) into a ‘direct’ and an ‘indirect’ component:⁴

$$\Delta \bar{p} = \overbrace{\frac{-c \text{Var}(p_i)}{\bar{w}}}^{\text{direct fitness effects}} + \overbrace{\frac{rb \text{Var}(p_i)}{\bar{w}}}^{\text{indirect fitness effects}} \quad (1)$$

Here, p_i denotes the individual gene frequency, with respect to the allele of interest, of the i^{th} individual in the population; and $\text{Var}(p_i)$ is the variance of p_i in the ancestral population.⁵ $\Delta \bar{p}$ is the change in the frequency of the allele between the ancestral and descendant populations, and $\bar{w}$ is the mean fitness in the ancestral population, where the ‘fitness’ of an individual is defined as the number of descendants it contributes to the descendant population. r is the coefficient of relatedness, defined, as in Queller (1992a), as a measure of the statistical association between the genotypes of social partners.⁶ c and b are, respectively, the coefficients of cost and benefit, defined, as in Queller (1992a), as partial regression coefficients in a regression model of fitness.

The key point to note about r , b and c is that they are population statistics, quantifying the overall associations in the population between the genotypes of social partners (r), one’s own genotype and one’s fitness (c) and one’s social partners’ genotypes and one’s fitness (b). With respect to r in particular, I should emphasize that it is a measure of genotypic assortment, which may or may not be attributable to kinship. As Hamilton (1975) stressed, there are sources of genotypic assortment that do not rely on kinship, including shared habitat preference, and so-called ‘greenbeard’ phenomena in which altruists (or bearers or some other social trait) recognize each other by means of a phenotypic marker (Dawkins, 1976; Gardner and West, 2010).

The above decomposition implies the following condition for positive change, which is a standard formulation of Hamilton’s rule:

$$\Delta \bar{p} > 0 \iff rb > c, \text{ provided } \text{Var}(p_i) \neq 0 \quad (\text{Hamilton’s rule})$$

Let us now compare this with a maximally general formulation of group selec-

⁴The same decomposition can be obtained for the change in a polygenic character, but in this case we should interpret p_i as a breeding value rather than an individual gene frequency (Queller, 1992a; Falconer and Mackay, 1996; Frank, 1998; Marshall, 2015). Here I focus on the single allele case.

⁵An individual’s gene frequency is the number of copies of the allele it possesses at the relevant locus in its genome, divided by its ploidy (Price, 1970). For a haploid organism, $p_i = 1$ if the individual has the allele and $p_i = 0$ otherwise.

⁶Formally, r is defined as the regression coefficient when we regress p_i on p'_i , where p'_i denotes the average p -value of the i^{th} individual’s social partners Grafen (1985); Queller (1992a).

tion theory: the multi-level Price equation (Price, 1972).⁷ This provides an alternative mathematical decomposition of change that applies whenever a population is group-structured. It partitions change into a ‘between-group’ and ‘within-group’ component:⁸

$$\Delta \bar{p} = \frac{\overbrace{\text{Cov}(W_k, P_k)}^{\text{between-group}}}{\bar{w}} + \frac{\overbrace{E_k[\text{Cov}_k(w_{jk}, p_{jk})]}^{\text{within-group}}}{\bar{w}} \quad (2)$$

Here, w_{jk} and p_{jk} denote the fitness and individual gene frequency (respectively) of the j^{th} member of the k^{th} group; while W_k and P_k denote (respectively) the mean fitness and group gene frequency (respectively) of the k^{th} group. $\text{Cov}(W_k, P_k)$, then, captures the covariance between a group’s gene frequency and its mean fitness, while $E_k[\text{Cov}_k(w_{jk}, p_{jk})]$ captures the average across groups of the within-group covariance between an individual’s gene frequency and its fitness.

This decomposition also implies a condition for positive change, which we might call ‘Price’s rule’:

$$\Delta \bar{p} > 0 \iff \text{Cov}(W_k, P_k) + E_k[\text{Cov}_k(w_{jk}, p_{jk})] > 0 \quad (\text{Price’s rule})$$

The argument for the ‘equivalence’ of the two conditions relies on the fact that both are derived from the Price equation (Price, 1970) with few additional assumptions. Both derivations assume that the allele is transmitted without bias. The only additional assumptions required for the derivation of Hamilton’s rule are that $\text{Var}(p_i) \neq 0$ and that p_i and p'_i are not collinear, so that the cost and benefit coefficients are well-defined. The only additional assumption required for the derivation of Price’s rule is that the population is group-structured. In all populations that satisfy the assumptions of both derivations (that is, in all group-structured populations in which alleles are transmitted without bias, p_i and p'_i are not collinear, and $\text{Var}(p_i) \neq 0$) both decompositions are correct and the following equivalence holds (Marshall, 2011; Frank, 2013; Birch and Okasha, 2015; Okasha,

⁷This way of formulating group selection is favoured by some theorists (such as Gardner and Grafen, 2009; Marshall, 2011; Gardner, 2015; Birch and Okasha, 2015; Okasha, 2016) and rejected by others in favour of an alternative approach known as contextual analysis (Heisler and Damuth, 1987; Damuth and Heisler, 1988; Goodnight *et al.*, 1992; Goodnight, 2013). Here I take the multi-level Price equation as my focal example of a multi-level approach; but, as I explain in the next subsection, the same general point could be made using contextual analysis as our focal example. The point of this section is simply that there is a close formal relationship between kin selectionist and multi-level methods, and the formal relationship is even closer if we take contextual analysis as our focal example of a multi-level method (see Okasha 2016, p. 440).

⁸The ‘Cov’ in $\text{Cov}(W_k, P_k)$ and the ‘E’ in $E_k[\text{Cov}_k(w_{jk}, p_{jk})]$ should be interpreted as ‘size-weighted’ functions in the sense of Price (1972); otherwise we need the further assumption that all groups are equal in size.

2016):

$$\Delta\bar{p} > 0 \iff rb > c \iff \text{Cov}(W_k, P_k) + E_k [\text{Cov}_k(w_{jk}, p_{jk})] > 0 \quad (3)$$

To understand the intuitive rationale for this, imagine the typical circumstances under which each condition would be satisfied for an altruistic trait controlled by a single gene. First, consider what is required for $rb > c$. It must be that bearers of the gene cluster together, so that the benefits of altruism fall differentially on bearers of the gene. Second, consider what is required for the selection against the trait within groups to be outweighed by selection for the trait between groups. Again, it must be that bearers of the gene cluster together, so that the heritable variation in fitness within groups is suppressed and the heritable variation in fitness between groups is boosted. Both approaches can thus be seen as alternative ways of capturing the fundamental requirement that bearers of the gene for altruism interact differentially with each other.

2.2 Limitations

The result in (3) is plainly an equivalence result of a sort. But we should be clear about what it does and does not show. Four main limitations are worth spelling out. Firstly, the result holds only in populations with a particular type of structure: group structure. Not all populations are group-structured, and the multi-level Price equation can be applied in the absence of group structure only by assigning organisms to groups arbitrarily, which deprives it of biological significance. I think the importance of this qualification has been understated in the literature (Godfrey-Smith 2006, 2008 is an important exception; see Section 4).

Secondly, the result involves comparing highly abstract, purely genetic formulations of kin selection and group selection theory, ignoring the complications that arise when we want to apply one of these approaches to analyse change in a particular ecological scenario. For example, in both the kin selectionist and multi-level modelling traditions, theorists tend to use phenotypic rather than genetic predictors of fitness when they have empirical applications in mind, since hypotheses about phenotypic selection gradients are easier to test empirically (Grafen, 1984; Queller, 1992b; Frank, 1998; Goodnight and Stevens, 1997; McGlothlin *et al.*, 2014). In fact, formal equivalence results can still be derived in relation to phenotypic versions of the two approaches, provided like is compared with like (Queller, 1992b; Birch and Marshall, 2014).

Thirdly, not all group selection theorists accept that the multi-level Price equation succeeds in separating the effects of selection at the group and individual levels. There is a prominent alternative—contextual analysis—advanced by Lorraine Heisler and John Damuth (Heisler and Damuth, 1987; Damuth and Heisler, 1988)

and by Charles Goodnight and colleagues (Goodnight *et al.*, 1992; Goodnight and Stevens, 1997; Goodnight, 2013). The clash between contextual analysis and the multi-level Price equation is long-running, and I do not aim to weigh into that debate here.⁹ It is sufficient to note that contextual analysis involves decomposing change in a very similar way to the generalized version of Hamilton’s rule: here too we decompose fitness using a regression model, but rather than taking the average genotype of the focal individual’s social partners as a predictor, we take properties of the focal individual’s social group (Heisler and Damuth, 1987). Thus the formal relationship between kin selectionist and multi-level methodologies would be even closer if we were to take contextual analysis as our flagship example of a multi-level approach, and it would become even clearer that the methodological differences between these traditions reflect divergent modelling preferences and explanatory interests rather than divergent empirical commitments (Okasha, 2016, p. 440).¹⁰

Fourthly, and most fundamentally, neither the generalized Hamilton’s rule nor Price’s rule, taken in isolation, says much about the causal processes driving evolutionary change. Like the Price equation itself, these rules are highly abstract, statistical results, compatible with a wide range of underlying causal explanations of change. There are many different causal explanations for the satisfaction of Hamilton’s rule (for reviews, see West *et al.*, 2007; Bourke, 2011). The same goes for Price’s rule: it provides a different way of carving up the change in gene frequency, but it too does so without implying anything in particular about the causes of change. It would therefore be a mistake to infer the identity of kin and group selection, conceived as causal processes responsible for change, from an equivalence result that merely concerns the relationship between two statistical conditions for change.

Of course, if we were to stipulate that by the term ‘kin selection’ we mean Hamilton’s rule and by the term ‘group selection’ we mean Price’s rule, evading the issue of causality, then there would be little to add to the equivalence result in (3). But I doubt whether this is the most useful way to employ these terms. After all, these terms intuitively refer to kinds of causal process—to things that actually happen in natural populations, and that feature in causal explanations—and not to formal methods, modelling traditions, or statistical conditions for change. I

⁹See Okasha (2006) for an overview of the issues.

¹⁰Goodnight (2013) suggests two ways in which the explanatory interests of the kin selectionist and contextual analysis traditions diverge: contextual analysis focusses on phenotypes whereas kin selectionist approaches tend to focus on genotypes (although many models in the kin selectionist tradition also focus on phenotypes, such as Queller, 1992b; Frank, 1998; McGlothlin *et al.*, 2014); and contextual analysis focusses on away-from-equilibrium change whereas kin selectionist approaches tend to focus on finding equilibria (although many models in the kin selectionist tradition also focus on away-from-equilibrium change, such as Hamilton, 1964; Queller, 1992b). These differences indicate different modelling preferences and explanatory interests, not deep disagreements about the causes of evolution.

think we should hold on to that intuition. I contend that the right moral to draw from the formal equivalence of Hamilton’s rule and Price’s rule, and other similar equivalence results, is not that kin and group selection are identical causal processes, but rather that purely statistical formalisms lack the resources to capture the causal distinction between them.¹¹

3 Individual- and Population-Centred Approaches

Broadly speaking, there are two approaches one can take to capturing the causal distinction. One is an individual-centred approach that explicates the distinction in terms of differences in the causal path that runs from a focal individual’s genes to its fitness. The other, which I will develop in this article, is a population-centred approach that explicates the distinction in terms of structural properties of populations. To provide a rationale for pursuing the second approach, I should comment briefly on the first.

Samir Okasha (2016) has recently pursued the first approach, drawing on the notion of a causal graph (see Spirtes *et al.*, 2000; Pearl, 2009). Okasha suggests that, in paradigm cases of group selection, a causal path runs ‘upwards’ from the individual gene frequency of a focal individual (p_i) to the local group’s gene frequency (P_i), then through the group gene frequency to the group mean fitness (W_i), and finally ‘downwards’ from group mean fitness and an ‘allocation mechanism’ to the focal individual’s fitness (w_i) (see Figures 8 and 9 in Okasha, 2016). In paradigm cases of kin selection, by contrast, there is no causal path running via the group means and no allocation mechanism. Instead, we have a causal path running directly, at the individual level, from the genes of one individual, via its own behaviour, to the fitness of another individual, and the fitness of the group is determined by the individual fitness of its members (see Figure 6 in Okasha, 2016). Figure 1, reprinted from Okasha (2016), depicts the relations between individual and group fitness in the two cases, without including genotypes (see Okasha, 2016, for more detailed figures).

Okasha’s graphs for paradigm cases of group selection posit ‘bottom-up’ causal relationships between individual gene frequencies and the group gene frequency and ‘top-down’ causal relationships between the group mean fitness and individual fitness values. I am uneasy with this aspect of the proposal, because I do not see the relationship between a set of individual properties and the group mean of those properties as one of causation. It strikes me as more accurate to describe this as a relationship of supervenience, because two groups cannot differ in their mean value of some property unless there is at least one difference between their respective sets of individual values.

¹¹A point also stressed by Birch and Okasha (2015); Okasha (2016), and Lehtonen (2016).




Figure 1: *An illustration of the causal graphs approach.* Case (a) is a paradigm case of kin selection; case (b) is a paradigm case of group selection. The mathematical relation between group fitness (W_i) and individual fitness (w_i) is the same in both cases: W_i is the group mean of w_i . However, in case (a) the individual fitness values metaphysically determine the group mean; whereas in case (b) the group mean, in conjunction with an allocation mechanism, metaphysically determines the individual fitness values (Figure reprinted from Okasha, 2016, Figure 4).

Foreseeing this problem, Okasha argues that, although these relationships would not normally be considered causal, they ‘can be depicted as if [they] were causal without violating the principles of causal modelling’ (Okasha, 2016, p. 449). However, this ‘as if’ move leads to a concern about whether the direction of these arrows is adequately constrained by the causal facts. In all cases, W_i supervenes on the set of w_i values. In some cases, Okasha draws a causal arrow from w_i to W_i ; in other cases, he draws a causal arrow from W_i to w_i . Why is the same relationship of supervenience to be represented in some cases by a top-down causal arrow and in other cases by a bottom-up causal arrow, and what constrains this choice?

Okasha’s view is that the choice depends on the ‘direction of metaphysical determination’ between w_i and W_i , which may be ascertained using ‘modal intuitions, empirical knowledge of the system being modelled, or both’ (p. 451). For example, Okasha suggests that, if w_i depends on the sharing of a group payoff, such as a large animal carcass, then it accords with intuition to say that the individual fitness values are metaphysically determined by the group mean fitness; whereas, if w_i depends only on payoffs obtained separately by individuals, such as smaller animals they have hunted individually, then it accords with intuition to say that the group mean fitness is metaphysically determined by the individual fitness values.

I do not share these intuitions: my intuition, for what it’s worth, is that a group

mean is always metaphysically determined by the individual values over which it averages, and that this remains the case even when the individual fitness values are causally explained by the sharing of a group resource. If this is right, then the arrow between w_i and W_i should always be a bottom-up arrow, and never a top-down arrow. However, I take it that intuitions on this question will differ, as will opinions regarding the evidential weight such intuitions merit. The deeper concern here is that Okasha's picture makes the classification of a process as one of kin or group selection dependent on such intuitions. This, I suggest, motivates the development of an alternative approach that can classify a process on the basis of its empirically observable features, without relying on intuitions about the direction of metaphysical determination that are subject to interpersonal variation.

A related but subtly different problem for Okasha's graphs concerns the arrow at the group level that runs from P_i to W_i . If we accept that the true relationship between a set of individual properties and the group mean of those properties is one of supervenience rather than causation, we run into traditional philosophical concerns about the causal efficacy of supervenient properties. Can group means cause other group means, or can the appearance of a causal relationship between two group means be explained away as a by-product (i.e. an 'epiphenomenon') of causation at the level of individual properties? So-called 'causal exclusion' arguments, a staple of the philosophy of mind for several decades, seem to have some purchase here (Kim, 2005).

This is not the place for a lengthy discussion of such arguments, or of the numerous responses to them (for a review of this area, see Robb and Heil, 2014). Okasha is right, I think, to set this issue to one side for his purposes. However, the way in which he does so leads to a problem. After acknowledging that the question of causal exclusion is a 'controversial metaphysical issue that is better not to prejudge' (p. 450), Okasha aims to sidestep the question by explicating the meaning of causal arrows between group variables in terms of hypothetical interventions on the supervenient property *and its lower-level supervenience base*:

[T]he following convention is adopted here: in a causal graph in which one variable supervenes on others, when we consider hypothetically intervening on the supervenient variable we do not hold fixed the variables on which it supervenes, but rather alter them to preserve consistency. Modulo this convention, causal arrows going out of supervenient variables, if any, can be understood in the usual way. (Okasha, 2016, p. 450)

Thus, we are to interpret a causal arrow running from P_i to W_i as implying that a counterfactual intervention on P_i *and the individual gene frequencies over which it averages* would lead to a change in W_i . This renders such arrows neutral regarding

the causal exclusion problem, since even a sceptic about full-blooded causation between group means should not object to the metaphysically thin relations of counterfactual dependence Okasha takes the group-level arrows in his graphs to imply. The trouble with this manoeuvre is that this thin sort of counterfactual dependence holds even in cases in which there is intuitively no group selection at work.

Consider G. C. Williams’s (1966) famous example of a population of deer, structured into herds. Herds containing faster deer are more successful than those containing slower deer, but only because faster individuals are more successful at evading predators than slower individuals. This is usually considered a case in which there is no genuine group selection—only covariance between a group’s mean running speed and its mean fitness caused by natural selection acting on individuals (Okasha, 2006, Ch. 3). Yet if one were to intervene on the mean running speed of a herd of deer, *altering the individual running speeds of the group members to preserve consistency*, this would make a difference to the group mean fitness. Okasha cannot consistently omit a causal arrow from P_i to W_i in such cases (as in Figure 5 in Okasha, 2016), given his apparent interpretation of the meaning of such an arrow.

I do not see these as fatal objections to the individual-centred approach Okasha pursues, but they are enough to motivate the development of an alternative. There are two key ideas at the heart of my approach that mark important departures from Okasha’s. First, I see the causal differences between kin selection and group selection as differences of degree, not all-or-nothing differences explicable in terms of the presence or absence of certain causal relationships. Second, I take it that the degree to which a selection process resembles a paradigm case of kin selection or group selection depends primarily on the structure of the population. Okasha’s graphs implicitly make assumptions about population structure (for example, a graph containing a ‘group gene frequency’ variable implicitly assumes the existence of groups), but they do not give population structure a central role. I favour an approach that explicitly accounts for the differences between kin and group selection in terms of the structural features of populations, bringing the role of population structure to the fore.

4 Two Influences: Hamilton and Godfrey-Smith

Before setting out the details of my proposal, I want to acknowledge (and highlight the insights of) two important influences. First, here are W. D. Hamilton’s (1975) own views on the relationship between kin and group (multi-level) selection:

If we insist that group selection is different from kin selection the term should be restricted to situations of assortment definitely not involving

kin. But it seems on the whole preferable to retain a more flexible use of terms; to use group selection when groups are clearly in evidence and to qualify with mention of ‘kin’ (as in the ‘kin group’ selection referred to by Brown), ‘relatedness’ or ‘low migration’ (which is often the cause of relatedness in groups), or else ‘assortation’, as appropriate. The term ‘kin selection’ appeals most where pedigrees tend to be unbounded and interwoven, as is so often the case with humans. (Hamilton, 1975, p. 337)

Although Hamilton in this passage initially sounds sceptical of there being a useful distinction between kin and group selection, he then proceeds to set out a nuanced way of thinking about that distinction. As I read it, he takes the view that the terminology of ‘kin selection’ and ‘group selection’ does track real and biologically important differences, but the differences that matter are differences of degree in aspects of population structure. The degree to which groups are ‘clearly in evidence’ matters, as does the degree to which assortment is explained by kinship. But the distinction is not clean or neat; it is not a dichotomy.

A second inspiration is Peter Godfrey-Smith’s (2006; 2008) work on the varieties of population structure. Godfrey-Smith contrasts group-structured populations with what he terms ‘neighbour-structured’ populations. In the former, social interactions are contained within sharply bounded, well-integrated groups in which everyone interacts with everyone else. In the latter, every individual interacts with its nearest neighbours, but there are no well-defined groups: there are only continuously overlapping networks centred on individuals. As Godfrey-Smith notes, one way to conceptualize the difference between these structures is in terms of the transitivity (or otherwise) of connections in social neighbourhoods. In the paradigm case of a group-structured network, the relation of fitness-affecting interaction is perfectly transitive (if A affects the fitness of B , and B affects the fitness of C , then A affects the fitness of C). By contrast, in the paradigm case of a neighbour-structured network—one in which each individual interacts with its four adjacent ‘von Neumann neighbours’ on a square lattice—the relation is perfectly intransitive: if A affects the fitness of B , and B affects the fitness of C , then A does *not* affect the fitness of C , assuming $A \neq C$. These should be seen as extreme cases: real social networks are likely to be neither perfectly transitive nor perfectly intransitive, but will instead have some intermediate level of transitivity.

The mathematical literature on network analysis gives us some formal tools with which to quantify the extent to which a network approximates these extreme cases. Network analysis has grown rapidly in recent years, and a great deal of work in this area has concentrated on the problem of identifying communities within networks (reviewed by Fortunato 2010). The starting point for any approach to this problem is to represent the whole-population social network as a graph

in which the individuals are the vertices (or nodes) and social interactions are the edges (or connections) between the vertices. Social neighbourhoods of focal individuals can then be represented as subgraphs. The vertices to which a focal vertex (v_i) is directly connected by an edge are known as its adjacent vertices. The subgraph $\mathcal{N}[v_i]$, comprising v_i and all vertices adjacent to v_i is known as the closed neighbourhood of v_i . This gives us the basic framework we need to start thinking more formally about the structure of social neighbourhoods.

One formal tool that is particularly useful for our purposes is the relative density of a social neighbourhood. Informally, the relative density compares the number of ‘inner’ connections joining up the members of a social neighbourhood to the number of ‘outer’ connections linking the members to other organisms outside the neighbourhood. Formally, the relative density is defined in terms of two other concepts: the internal and total degree of a subgraph. Consider the subgraph $\mathcal{N}[v_i]$, representing the closed neighbourhood of a focal vertex. The internal degree of an arbitrary vertex v_j in $\mathcal{N}[v_i]$ is the number of edges directly linking v_j to other vertices within $\mathcal{N}[v_i]$; the external degree of v_j is the number of edges directly linking it to vertices outside $\mathcal{N}[v_i]$; and the total degree of v_j is the sum of its internal and external degrees. The internal degree of the subgraph $\mathcal{N}[v_i]$ is then defined as the sum of the internal degrees of its vertices, and the external and total degrees of the subgraph are likewise defined as the sum of the external and total degrees (respectively) of its vertices. The relative density of $\mathcal{N}[v_i]$ is the ratio of its internal degree to its total degree (Fortunato, 2010, p. 85).

In Godfrey-Smith’s paradigmatic neighbour-structured population, in which each organism interacts with its four von Neumann neighbours on a square lattice, the subgraph defined by a focal individual and its von Neumann neighbourhood has a relative density of $2/5$ (0.4): the internal degree is 8 and the total degree is 20. By contrast, the subgraph defined by an hermetically sealed social group, with no outward connections, has a relative density of 1. An intermediate case is a square lattice in which each individual interacts with its eight ‘Moore neighbours’ (including, in addition to its four von Neumann neighbours on each side, the four neighbours on the corners between these sides): the internal degree is 40 and the total degree is 72, implying a relative density of $5/9$ (0.55).

The relative density is very sensitive to isolation—a subgraph with some internal connections and no outward connections will always have a relative density of 1, no matter how poorly integrated it is—but, for any subgraph with some outward connections, it is also sensitive to the extent to which the subgraph is internally joined-up. For these reasons, the relative density provides an attractive way of quantifying the extent to which well-defined groups are ‘clearly in evidence’ in a population.

5 K and G

My proposal is that we conceptualize the distinction between kin and group selection in terms of gradated differences in two key structural properties of populations. I will label these properties as K (for ‘kin-structure’) and G (for ‘group-structure’).¹³ Kin selection, roughly speaking, is selection on indirect fitness differences ($rb \neq 0$) that occurs in a high- K population (a population with a high degree of kin-structure); whereas group selection, roughly speaking, is selection on indirect fitness differences ($rb \neq 0$) that occurs in a high- G population (a population with a high degree of group-structure).

To be clear, this proposal is not intended to capture all current usages of the terms ‘kin selection’ and ‘group selection’. Rather, it is a proposal about how these concepts should be used, if we want them to mark a real and evolutionarily significant distinction among selection processes. It is to some extent a revisionary proposal, although, as I have been emphasizing, I see it as well-aligned with Hamilton’s own views on how the distinction should be drawn.

Before explaining K and G , let me stress the condition that, for either kin or group selection to occur, it must be the case that $rb \neq 0$ in the population as a whole. In other words, both kin and group selection act on indirect fitness differences, and therefore rely on the presence of positive genotypic assortment. The requirement may be less restrictive than it initially appears, since, as I emphasized in Section 2, positive genotypic assortment can arise from sources that do not rely on genealogical kinship, such as shared habitat preference and greenbeard effects. If $rb = 0$, then the selection process at work relies on direct fitness effects alone, and I claim that to count such a process as one of kin or group selection unhelpfully obscures this fact. If what is on offer is a direct fitness explanation, we should not invoke these concepts. While I hope this sounds reasonable on first hearing, many group selection theorists allow that group selection can occur when $rb = 0$, so I will comment further on this issue below (in Section 6).

¹²But note that, if a group is isolated but poorly connected internally, this will tend to show up in lower relative densities for the social neighbourhoods of individuals at the periphery of the group. For example, imagine a subgraph comprising a central individual and its four von Neumann neighbours that has no outward connections to the rest of the population. The central individual’s social neighbourhood has a relative density of 1, but, for the four peripheral individuals, the relative density of their social neighbourhood is 0.4. For this reason, the average relative density over all neighbourhoods will tend to be quite low when groups are poorly integrated internally, even if the groups are isolated from each other.

¹³This form of labelling is inspired by that of Godfrey-Smith (2009).

5.1 K

K , the degree of kin-structure in a population, is intended to capture the overall extent to which genealogical relatives interact differentially with respect to the character of interest. Accordingly, I will refer to populations in which there is a high degree of differential interaction between relatives as ‘high- K ’ populations; and I will refer to populations in which there is no tendency for relatives to interact differentially as ‘zero- K ’ populations.

I do not intend to commit to a single quantitative measure of K , firstly because I want to allow that different measures may be appropriate in different contexts, and secondly because I do not need to commit to a measure in order to use K to make qualitative comparisons among populations (cf. Godfrey-Smith, 2009; Queller and Strassmann, 2009). However, for the purpose of fixing ideas, it may be helpful to think of K as the correlation between social partner genotypes (with respect to the character of interest) that would obtain in the absence of any kinship-independent sources of such correlation, such as greenbeard effects.¹⁴

How high does the degree of differential interaction between relatives have to be before we have a case of kin selection? Because we are dealing with a continuum of cases here, any cut-off will be a pragmatic choice, and it is arguably best to avoid any such cut-off. Following Godfrey-Smith (2009), I prefer to talk of ‘marginal’ and ‘paradigm’ cases. Paradigm cases of kin selection occur in high- K populations. When we have non-zero rb but very low K , either because r is very low or because it is largely generated by kinship-independent mechanisms, we have at best a marginal case of kin selection, and such a selection process is probably more aptly described in other terms. Human evolution may be an example of a marginal case, since estimates based on studies of modern hunter-gatherers suggest a value of genetic relatedness of around 0.05 in such societies (Hill *et al.*, 2011; Bowles and Gintis, 2011).

One might ask: why does K matter? Why is this a structural property worth estimating? Why are comparisons among populations, in regard to their degree of K , worth making? My answer is that kin-structure has a special role to play in generating the conditions for the evolution of stable altruistic or spiteful behaviour. Genetic correlations can certainly arise without kinship, as shown by greenbeard phenomena (Hamilton, 1975; Dawkins, 1976). But there is a standard concern regarding greenbeard effects: altruism that relies on this mechanism will be stable only if, for some reason, the expression of the altruistic behaviour cannot be suppressed without also suppressing the phenotypic marker (in Dawkins’s famous example, a literal green beard) that attracts benefits from others. If this selective suppression is possible, then it will pay to be a ‘falsebeard’: an organism

¹⁴In cases where kinship-independent sources are actually absent, this is related to r but not identical to it, since r is a regression coefficient rather than a correlation coefficient.

who expresses the marker without expressing the altruism. By contrast, genetic correlations generated by kinship-dependent mechanisms are not so easy to subvert, because kinship generates genetic correlation at every locus in the genome (Ridley and Grafen, 1981; Okasha, 2002; Gardner and West, 2010).

5.2 G

G , the degree of group-structure in a population, is intended to capture the overall extent to which a population contains well-defined social groups, at the right grain of analysis for generating non-zero rb , that are stable over the course of the life-cycle. A ‘high- G ’ population is one in which groups are well-integrated, highly stable and effectively insulated from other groups, with no room for ambiguity regarding group membership. John Maynard Smith’s (1964) haystacks model, in which we imagine social interaction and reproduction occurring in isolated sub-populations (envisioned as haystacks inhabited by mice), with occasional mixing events, is a good example of this.¹⁵ A ‘low- G ’ population is one in which, although interaction is locally structured to some extent, there are no discrete, well-defined social groups to speak of, because—as in the von Neumann neighbour-structured populations of Godfrey-Smith (2008)—social neighbourhoods blur continuously into one another. A ‘zero- G ’ population is one in which we do not even have neighbour-structure: individuals interact with social partners drawn from the population as a whole, with no regard to their spatial location.

The qualification ‘at the right grain of analysis for generating non-zero rb ’ merits emphasis. For example, one might worry that all populations of multicellular animals are ultimately high- G populations: after all, there is always group structure if one looks at a fine enough grain of analysis, because one can always describe individual animals as groups of cells. However, this sort of ‘group-structure’ is at the wrong grain of analysis if we want to explain the evolution of an organism-level social phenotype manifested in interactions between organisms. The right grain of analysis is that of the organism-level social network defined by fitness-affecting interactions with respect to the phenotypic character of interest. The population is ‘high- G ’ if that network can be subdivided into sharp and stable social groups.

As with K , I do not intend to commit to a single quantitative measure of G . I suspect there is no perfect measure, and that the most appropriate measure will depend on the context, because the relative importance of the different properties that contribute to G —internal integration, external isolation, and stability over time—will depend on the context. But again, for the purpose of fixing ideas, it may be helpful to have a possible measure in mind. One possible measure with

¹⁵The place of a haystacks model on the K -axis depends on the parameter values: the size of the founding population, the assortativity of group formation and the time of isolation.

attractive features is the average, taken over all individuals in the population and over an appropriate time period, of the relative density of a focal individual's social neighbourhood. As we saw in Section 4, this ranges between 0 and 1, and places von Neumann neighbour-structure at 0.4, Moore neighbour-structure at 0.55, and perfectly integrated, hermetically sealed groups at 1. The range 0 to 0.4 is occupied by social structures in which the average social neighbourhood has a greater external degree than we see in a von Neumann neighbour-structured population, without displaying significantly more internal integration.¹⁶ The range 0.55 to 1 is occupied by social structures that display less internal integration and/or external isolation than in the idealized extreme case, but more internal integration and/or external isolation than we see in a Moore neighbourhood.

The relative density is well suited to measuring internal integration and external isolation of social neighbourhoods, but less well suited to measuring their stability, underlining the point that there is probably no single perfect measure of G . However, the time-average of the relative density over an extended time period will convey something about the stability of groups over that time period: if well-defined groups are ephemeral and dissolve soon after forming, the relative density will be high while they exist but lower once they have dissolved, resulting in a lower time-average than in a population with more stable group-structure. So, while the relative density is not intended as a measure of group stability, time-averages of the relative density may sometimes be useful for that purpose.

As with K , one might ask: why does G matter? Why is this a structural property worth measuring? Why are comparisons among populations, in regard to their degree of G , worth making? My answer to this question is to point to the special role of high- G populations in evolutionary transitions in individuality. A population that is high- G contains identifiable, stable, bounded, higher-level entities—namely, social groups—formed of collections of lower-level entities. These groups are not automatically higher-level individuals. I take it that higher-level individuality requires some process of collective reproduction (see Godfrey-Smith, 2009), as well as the presence of mechanisms that suppress selection within (or, in Godfrey-Smith's memorable terminology, 'de-Darwinize') the groups (Michod, 1999; Godfrey-Smith, 2009; Queller and Strassmann, 2009; Clarke, 2013). Nevertheless, group-structure is clearly an important precondition for the evolution of higher-level individuals. When we identify a population as high- G , we cannot conclude that a transition is underway, but we can conclude that an important precondition for such a transition has been met.

¹⁶For example, imagine a structure in which every individual has five social partners drawn at random from a very large population. This is likely to result in social neighbourhoods with relative densities of around $1/3$.

5.3 K - G space

K and G can be imagined as the axes of a two-dimensional space, and we can think of kin selection and group selection as large, overlapping regions of that space. Paradigm cases of kin selection occur in high- K populations: they are cases in which we find selection on indirect fitness differences in a population with a fairly high degree of relatedness between social partners, and with kinship-dependent mechanisms serving as the main source of this relatedness. Paradigm instances of group selection occur in high- G populations: they are cases in which we find selection on indirect fitness differences in a population in which social interaction is structured by stable, well integrated and sharply bounded groups. The distinction is not sharp, but nor is it merely arbitrary or conventional.




Figure 2: K - G space. Kin selection and group selection can be conceptualized as overlapping regions of a 2D space defined by the variables K and G . Locations of points are approximate and for illustration only (see the main text for commentary on some of the points).

Figure 2 provides a visualization of K - G space, illustrated with some notable cases. The placement of the points is not exact and is open to debate: the aim is simply to provide an intuitive visualization. In the bottom-left corner, we have populations that are low- K and low- G —populations with neither kin-structure nor group-structure. An important class of examples are so-called ‘well-mixed’ populations in which individuals interact at random such that no pair of individuals is any more likely to interact than any other. These are populations in which neither kin nor group selection can be said to occur, since the required structural features are entirely lacking.

As we move up the K -axis, we come to populations in which organisms still interact with sets of individuals drawn from the whole population with no regard to their spatial location, rather than interacting in structured local neighbourhoods, but in which there is some greater-than-chance probability of interacting with a relative (e.g. a sibling). Models of this sort have a long history in social evolution theory and continue to be studied (e.g. Queller, 1984; Allen and Nowak, 2015). These are aptly described as cases of kin selection, but, since interactions are not contained within localized social groups, they are not aptly described as cases of group selection.

As we move along the G -axis, we come to Godfrey-Smith’s neighbour-structured populations, in which there are discernible local neighbourhoods that structure interaction, but nothing yet resembling well-defined social groups. A square lattice in which organisms interact with their von Neumann neighbours and are assigned to vertices at random, with no limited dispersal, is a zero- K version of this. As we go up the K -axis here, introducing differential interaction between kin due to limited dispersal, we arrive at populations that are high- K but still fairly low- G . Models of so-called viscous populations that make use of von Neumann neighbourhoods and similar structures, such as the models of D. S. Wilson *et al.* (1992), belong in this area; their precise position will depend on the parameter values. In models of haploid organisms, very high levels of relatedness can be attained due to limited dispersal from the birth site (in one of the simulations discussed by Wilson *et al.* 1992, $r = 0.59$), and this is reflected in the figure in the value of K . Moving further along the G -axis, we arrive at the discrete ‘trait-groups’ of D. S. Wilson (1975), which are externally isolated and fairly well integrated (at least with respect to the trait of interest) while they exist, but which are not stable for long periods and are typically outlived by their members. These can be anywhere on the K -axis depending on the role of kinship in generating assortative grouping, but for illustrative purposes I have put them at low- K .

In the bottom-right corner, there are populations that are low- K and high- G . Here, groups are ‘clearly in evidence’ but groups are not composed of close kin. Bowles and Gintis’s (2011) models of human evolution belong in this region.

Bowles and Gintis assume that early human populations were structured into well-defined, stable groups with low relatedness. Finally, as we go up the K -axis to the top-right, we arrive at populations that are high in both K and G . These are the cases for which Hamilton favoured the term ‘kin-group selection’. In these populations, there is sharp and stable group structure *and* a high degree of genetic correlation between social partners due to kinship-dependent mechanisms.

The evolution of multicellularity is a source of extreme cases in this corner (Maynard Smith and Szathmáry, 1995; Michod, 1999). Consider colonial algae such as *Gonium*: the colonies are clonal, implying high K , and the group structure is sharp and stable over the course of the life-cycle, implying high G . There is little to be gained by arguing over whether the selection processes that operate in these populations are cases of kin or group selection, because they have the core structural features of both. The term ‘kin-group selection’, which removes the misleading appearance of competing hypotheses, seems apt for processes in this region of the space.

6 The $rb \neq 0$ Requirement

The proposal I have advanced includes the requirement that, for either kin or group selection to occur, it must be the case that $rb \neq 0$. I noted above that this requirement, traditionally associated with kin selection, might prove controversial as a requirement on group selection, and it is now time to elaborate further on the consequences of, and justification for, this requirement.

Let us first consider some of its implications. One is that not all processes of natural selection occurring in populations that contain groups will qualify as cases of group selection. Consider again G. C. Williams’s scenario in which a group containing fast-running deer outperforms a group containing slower-running deer because the faster deer, as individuals, evade predators more easily. This is not group selection on my account, assuming the trait’s advantage arises entirely from its direct fitness effects (Okasha, 2006). The intuitive motivation for excluding these cases is that, although a form of group-structure is present, it plays no role in explaining the selection for fast running. Similarly, processes of natural selection that involve interactions among relatives do not qualify as kin selection if the interactions fail to generate non-zero rb , perhaps because the interactions are not fitness-affecting, or because social partners, though related, are not differentially related relative to the population average.

Awkward cases arise when, although $rb = 0$, intergroup conflict plays an essential role in the generation of a direct fitness benefit. Sterelny’s (2013) hypothesis regarding the evolution of hierarchy in early Holocene human societies provides an interesting example. Sterelny proposes that ruling elites were tolerated by the

majority because the intense and frequent intergroup warfare of the early Holocene put a fitness premium on strong and centralized military leadership. Groups with strong leaders were more successful in warfare, causing traits associated with hierarchy to spread. Is this a group selection hypothesis? In one sense it is, because the advantage of acquiescing to the demands of an elite depends on the existence of intergroup conflict. But on my proposal, it is not, because the explanation on offer is a direct fitness explanation. Norms of acquiescence evolve because, at an individual level, it pays in direct fitness terms to adopt them; there is no requirement here that $rb > 0$. Not a lot hangs on how we classify these cases, and I propose that we resist the urge to describe them as cases of group selection. In so doing, however, we should take care not to forget that direct fitness explanations can still appeal to intergroup conflict as a source of direct fitness benefit.

If $rb \neq 0$, must we conclude that one or other of kin and group selection is at work? Not necessarily, for recall that the requirement is intended as a necessary but not sufficient condition. Consider greenbeard effects. Populations in which the only genetic correlations between social partners are owed to greenbeard phenomena belong in the bottom-left quadrant of Figure 2. They do not require group-structure, provided we assume that bearers of the greenbeard marker can still seek each other out successfully in a non-group-structured population, but nor do they rely on differential interaction between genealogical kin, since bearers of the greenbeard gene need not be kin in this sense.

Gardner *et al.* (2011) regard greenbeard effects as a form of kin selection, broadly construed. I would say that these are, at most, marginal cases. It is important to distinguish clearly between cases in which genealogical kinship is pivotal, as in paradigm cases of kin selection, and marginal cases in which kinship-dependent mechanisms are a minor or negligible contributor to r . In microbes, gene mobility provides another source of locus-specific genetic correlation that does not rely directly on identity of alleles by virtue of descent, as noted by Rankin *et al.* (2011), Mc Ginty *et al.* (2013) and Birch (2014). Rankin *et al.* (2011) suggest that this too can be regarded as a form of kin selection, but I regard it as a highly marginal case. Note that, in some cases, kinship-independent mechanisms may operate in conjunction with kinship-dependent mechanisms, such that both contribute to the value of r . We therefore have a continuum here—not a dichotomy—ranging from highly marginal cases of kin selection in which shared ancestry is wholly unimportant to paradigm cases in which it is essential.

The main reason I anticipate resistance to the idea that group selection requires $rb \neq 0$ is that $rb \neq 0$ is neither necessary nor sufficient for there to be variation in fitness between groups. Non-zero rb is unnecessary because, as in the aforementioned cases of Williams’s fast-running deer and Sterelny’s explanation for acquiescence to hierarchies, there can be fitness variation between groups even

though direct fitness effects fully account for this variation. However, as noted above, I think it is unhelpful to classify these as cases of group selection. Non-zero rb is insufficient due to the possibility of soft selection with local population regulation, as discussed by Heisler and Damuth (1987); Goodnight *et al.* (1992) and Okasha (2006). In such cases, we have a group-structured population, but each group makes the same, fixed contribution to the next generation, and all fitness variation occurs within groups. Yet grouping is assortative—altruists interact differentially with other altruists—leading to non-zero rb . This population is high- G , suggesting a paradigm case of group selection by the lights of my account, but there is no variation in fitness between groups.

If one takes variation in fitness between groups to be the mark of group selection, then one should take $\text{Cov}(W_k, P_k) \neq 0$, not $rb \neq 0$, as the minimal statistical requirement all cases of group selection must satisfy. This would add an extra layer of complexity to the account, since kin and group selection would then differ in their minimal statistical requirements as well as in their commitments regarding population structure. However, I resist this amendment, because I do not see a compelling case for regarding variation in group fitness as necessary for group selection. If well-defined group-structure is implicated in generating non-zero rb , I take the view that the selection process can be aptly described as one of group selection, even if groups do not vary in mean fitness.¹⁷

Why insist that group selection must require fitness variation between groups? I see two main motivations. One is a desire that the conditions for group selection should be directly analogous to the conditions for natural selection in a population of individuals, but with ‘groups’ substituted for ‘individuals’. The conditions for natural selection include fitness variation among individuals, so group selection must require fitness variation among groups for a direct analogy to hold.

I reply that, although a direct analogy between the conditions for individual and group selection would be elegant, it does not deserve high priority. It is helpful here to invoke Heisler and Damuth’s (1987) MLS1/MLS2 distinction, and to recall that the type of group selection at issue in this article is the MLS1 type. In MLS2, groups are higher-level individuals or proto-individuals, reproducing in their own right. In this context, a direct analogy between the conditions for individual and group selection seems important, because ‘group selection’ in this sense is simply

¹⁷My position here is well-aligned with what Okasha (2006) calls the ‘neighbour approach’ of Nunney (1985), which diagnoses group selection in a group-structured population whenever there is positive rb . This is closely related to, although not identical to, the ‘contextual approach’ of Heisler and Damuth (1987). Both approaches decompose change using regression models of fitness. The difference is that the contextual approach uses regression models that take group characters (e.g. group gene frequencies) as predictors, whereas the neighbour approach uses neighbourhood characters (e.g. the average gene frequency of the focal individual’s social partners). The latter corresponds to the partition of change represented in Hamilton’s rule.

a higher-level form of individual selection. In MLS1, by contrast, groups structure interaction at the lower-level but do not reproduce in their own right, making the need for a direct analogy seem less pressing (Okasha, 2006). The priority, in my view, is that the category of group selection demarcates (albeit not sharply) a real and evolutionarily significant class of selection processes—a class that is worth distinguishing from the class of kin selection processes. On my account, it does this: the distinction between kin and group selection highlights real and evolutionarily significant differences in population structure within the class of indirect fitness explanations of change.

A second motivation is that group selection should be apt to generate group adaptation, and there can be no group adaptation without fitness variation between groups (Gardner and Grafen, 2009). I reply that, although the connection between group selection and group adaptation is important, especially in the context of evolutionary transitions, there should be no requirement that group selection must be apt to generate group adaptation in all cases. It is enough that this can happen under some further conditions—conditions that will include variation in fitness between groups. To insist that group selection must require variation in fitness between groups is, I think, to insist on too close a link between group selection and group adaptation.

7 Levels of Organization

One final clarifying remark deserves special emphasis: both kin-structured (high- K) and group-structured (high- G) populations can occur at multiple levels of biological organization. If we take a group-structured population of base-level entities and ‘frameshift’ up a level to consider the population of groups, this higher-level population will itself have a position in K - G space.

The population of groups may have higher-order group-structure: there may be sharply bounded meta-groups, or groups of groups, defined by patterns of social interaction. This may lead to higher-level group selection. For example, there might be higher-level group selection for genes that promote cooperation among members of the same meta-group who are located in different first-order groups.

The population of groups may also be kin-structured. This will be the case if groups form well-defined lineages, and if groups that are closely related genealogically are more likely to interact with each other than groups that are not. This may lead to a higher-level form of kin selection. For example, if nearby groups tend to be ‘offspring’ groups of the same ‘parent’ group, there might be higher-level kin selection for genes that promote cooperation among nearby groups. This may occur even if the population of groups is simply a viscous population, with no well-defined meta-groups.

These ideas may sound strange at first hearing, but they are simply unusual ways of describing something familiar. From the point of view of social evolution theory, multicellular organisms can be regarded as particularly well-integrated social groups of cells (Queller and Strassmann, 2009; Bourke, 2011). From this perspective, populations of animals are populations of groups of lower-level entities, and standard cases of kin selection and group selection occur in such populations.

This should alert us to the possibility of kin and group selection occurring at higher levels of biological organization than we usually envisage: that is, in kin- or group-structured populations of groups of organisms. For example, in many ant species we find ‘supercolonies’, each consisting of multiple distinct nests. This leads to the idea that supercolonies may be created and maintained by group selection acting on groups of nests—a possibility highlighted by Bourke (2011) and investigated by Kennedy *et al.* (2014). Moreover, dispersal of offspring nests from parental nests is limited within supercolonies, raising the possibility that, within a supercolony, kin selection at the level of the nest favours cooperation between adjacent nests—a possibility discussed by Chapuisat *et al.* (1997).

8 The Key Substantive Questions

Debates surrounding kin and group selection are easily derailed by semantic confusion. This, combined with the plethora of ‘equivalence results’ described in Section 2, gives rise to the suspicion that there are no worthwhile debates to be had here at all: that all the disputes on these issues are merely verbal disputes. But I think this suspicion is misplaced. By identifying kin and group selection with overlapping regions of K - G space, we make room for worthwhile debates about the resemblance of a given selection process (such as early human evolution, or social evolution in microbes, or the evolution of eusociality in insects) to a paradigm case of kin or group selection. A population’s position in K - G space will depend on the answers to the following questions:

1. How high is K in the population? That is, how strong are the genetic correlations between social partners, and how important are kinship-dependent sources, as opposed to greenbeard effects and other kinship-independent sources, in generating those correlations?
2. How high is G in the population? That is, how internally integrated, sharply bounded and stable is the group-structure at the relevant grain of analysis?

These are substantive, empirical questions: questions it takes empirical inquiry, and not just stable semantic conventions, to settle. Moreover, the position of a population in K - G space has significant consequences for its evolutionary fate: K

makes a difference to the long-term stability of altruism in the population, while G makes a difference to its chances of undergoing an evolutionary transition in individuality. So, although the distinction between kin and group selection is not sharp, these concepts can still provide a useful way of framing meaningful debates about the importance of kin-structure and group-structure in real processes of social evolution. Kin and group selection correspond to large, overlapping regions of a space of population structures, and real populations can be found throughout these regions.

Acknowledgements

I thank Caroline Birch, Andrew Buskell, Ellen Clarke, Peter Godfrey-Smith, Heikki Helanterä, Jussi Lehtonen, James Marshall, Samir Okasha, Jun Otsuka, Tobias Uller, the UCL Human Evolutionary Ecology Group, reading groups at UC Irvine and the Australian National University, three anonymous reviewers, and an associate editor at this journal for their advice and feedback. I also thank audiences at UC Irvine, the LSE, the University of Utah, the University of Cambridge, and Philosophy of Biology at Dolphin Beach for their comments and questions. This work was supported by a Philip Leverhulme Prize from the Leverhulme Trust.

*Department of Philosophy, Logic and Scientific Method
The London School of Economics and Political Science
Houghton Street, London, WC2A 2AE, UK
j.birch2@lse.ac.uk*

Bibliography

- Allen, B. and Nowak, M. A. [2015]: ‘Games among Relatives Revisited’, *Journal of Theoretical Biology*, **378**, pp. 103–116.
- Birch, J. [2014]: ‘Gene Mobility and the Concept of Relatedness’, *Biology and Philosophy*, **29**, pp. 445–476.
- Birch, J. and Marshall, J. A. R. [2014]: ‘Queller’s Separation Condition Explained and Defended’, *American Naturalist*, **184**(4), pp. 531–540.
- Birch, J. and Okasha, S. [2015]: ‘Kin Selection and Its Critics’, *BioScience*, **65**(1), pp. 22–32.
- Bourke, A. F. G. [2011]: *Principles of Social Evolution*, Oxford: Oxford University Press.

- Bowles, S. and Gintis, H. [2011]: *A Cooperative Species: Human Reciprocity and its Evolution*, Princeton: Princeton University Press.
- Chapuisat, M., Goudet, J. and Keller, L. [1997]: ‘Microsatellites Reveal High Population Viscosity and Limited Dispersal in the Ant *Formica paralugubris*’, *Evolution*, **51**, pp. 475–482.
- Clarke, E. [2013]: ‘The Multiple Realizability of Biological Individuals’, *Journal of Philosophy*, (8), pp. 413–435.
- Damuth, J. and Heisler, I. L. [1988]: ‘Alternative Formulations of Multilevel Selection’, *Biology and Philosophy*, **3**, pp. 407–430.
- Dawkins, R. [1976]: *The Selfish Gene*, New York: W. W. Norton and Company.
- Dawkins, R. [1982]: *The Extended Phenotype: The Gene as the Unit of Selection*, New York: W. W. Norton and Company.
- Falconer, D. S. and Mackay, T. F. C. [1996]: *Introduction to Quantitative Genetics*, London: Longman, 4th edition.
- Fortunato, S. [2010]: ‘Community Detection in Graphs’, *Physics Reports*, **486**, pp. 75–174.
- Frank, S. A. [1998]: *Foundations of Social Evolution*, Princeton, NJ: Princeton University Press.
- Frank, S. A. [2013]: ‘Natural Selection. VII. History and Interpretation of Kin Selection Theory’, *Journal of Evolutionary Biology*, **26**, pp. 1151–1184.
- Gardner, A. [2015]: ‘The Genetical Theory of Multilevel Selection’, *Journal of Evolutionary Biology*, **28**, pp. 305–319.
- Gardner, A. and Grafen, A. [2009]: ‘Capturing the Superorganism: A Formal Theory of Group Adaptation’, *Journal of Evolutionary Biology*, **22**, pp. 659–671.
- Gardner, A. and West, S. A. [2010]: ‘Greenbeards’, *Evolution*, **64**, pp. 25–38.
- Gardner, A., West, S. A. and Wild, G. [2011]: ‘The Genetical Theory of Kin Selection’, *Journal of Evolutionary Biology*, **24**, pp. 1020–1043.
- Godfrey-Smith, P. [2006]: ‘Local Interaction, Multilevel Selection, and Evolutionary Transitions’, *Biological Theory*, **1**, pp. 372–380.

- Godfrey-Smith, P. [2008]: ‘Varieties of Population Structure and the Levels of Selection’, *British Journal for the Philosophy of Science*, **59**, pp. 25–50.
- Godfrey-Smith, P. [2009]: *Darwinian Populations and Natural Selection*, Oxford: Oxford University Press.
- Goodnight, C. J. [2013]: ‘On Multilevel Selection and Kin Selection: Contextual Analysis Meets Direct Fitness’, *Evolution*, **67**(6), pp. 1539–1548.
- Goodnight, C. J., Schwartz, J. M. and Stevens, L. [1992]: ‘Contextual Analysis of Models of Group Selection, Soft Selection, Hard Selection, and the Evolution of Altruism’, *American Naturalist*, **140**, pp. 743–761.
- Goodnight, C. J. and Stevens, L. [1997]: ‘Experimental Studies of Group Selection: What Do They Tell Us About Group Selection in Nature?’, *American Naturalist*, **150**(S1), pp. s59–s79.
- Grafen, A. [1984]: ‘Natural Selection, Kin Selection and Group Selection’, in J. R. Krebs and N. B. Davies (eds), *Behavioural Ecology*, Oxford: Blackwell, 2nd edition, pp. 62–84.
- Grafen, A. [1985]: ‘A Geometrical View of Relatedness’, *Oxford Surveys in Evolutionary Biology*, **2**, pp. 28–89.
- Hamilton, W. D. [1964]: ‘The Genetical Evolution of Social Behaviour I and II’, *Journal of Theoretical Biology*, **7**, pp. 1–52.
- Hamilton, W. D. [1975]: ‘Innate Social Aptitudes of Man: An Approach from Evolutionary Genetics’, in R. Fox (ed.), *Biosocial Anthropology*, New York: Wiley, pp. 133–55.
- Heisler, I. L. and Damuth, J. [1987]: ‘A Method for Analyzing Selection in Hierarchically Structured Populations’, *American Naturalist*, **130**, pp. 582–602.
- Hill, K. R., Walker, R. S., Božičević, M., Eder, J., Headland, T., Hewlett, B., Hurtado, A. M., Marlowe, F., Wiessner, P. and Wood, B. [2011]: ‘Co-Residence Patterns in Hunter-Gatherer Societies Show Unique Human Social Structure’, *Science*, **331**(6022), pp. 1286–1289.
- Kennedy, P., Uller, T. and Helanterä, H. [2014]: ‘Are ant supercolonies crucibles of a new major transition in evolution?’, *Journal of Evolutionary Biology*, **27**(9), pp. 1784–1796.
- Kim, J. [2005]: *Physicalism, or Something Near Enough*, Princeton, NJ: Princeton University Press.

- Lehtonen, J. [2016]: ‘Multilevel Selection in Kin Selection Language’, *Trends in Ecology and Evolution*, **31**(10), pp. 752–762.
- Marshall, J. A. R. [2011]: ‘Group Selection and Kin Selection: Formally Equivalent Approaches’, *Trends in Ecology and Evolution*, **26**, pp. 325–332.
- Marshall, J. A. R. [2015]: *Social Evolution and Inclusive Fitness Theory: An Introduction*, Princeton, NJ: Princeton University Press.
- Maynard Smith, J. [1964]: ‘Group Selection and Kin Selection’, *Nature*, **200**, pp. 1145–1147.
- Maynard Smith, J. [1976]: ‘Group Selection’, *Quarterly Review of Biology*, **21**, pp. 20–29.
- Maynard Smith, J. and Szathmáry, E. [1995]: *The Major Transitions in Evolution*, Oxford: Oxford University Press.
- Mc Ginty, S. E., Lehmann, L., Brown, S. P. and Rankin, D. J. [2013]: ‘The Interplay between Relatedness and Horizontal Gene Transfer Drives the Evolution of Plasmid-Carried Public Goods’, *Proceedings of the Royal Society of London B: Biological Sciences*, **280**, pp. 20130400.
- McGlothlin, J. W., Wolf, J. B., Brodie, E. D. and Moore, A. J. [2014]: ‘Quantitative Genetic Versions of Hamilton’s Rule with Empirical Applications’, *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, **369**(1642), pp. 20130358.
- Michod, R. E. [1999]: *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*, Princeton, NJ: Princeton University Press.
- Nowak, M. A., Tarnita, C. E. and Wilson, E. O. [2010]: ‘The Evolution of Eusociality’, *Nature*, **466**, pp. 1057–1062.
- Nunney, L. [1985]: ‘Group Selection, Altruism, and Structured-Deme Models’, *American Naturalist*, **126**, pp. 212–230.
- Okasha, S. [2002]: ‘Genetic Relatedness and the Evolution of Altruism’, *Philosophy of Science*, **69**, pp. 139–149.
- Okasha, S. [2006]: *Evolution and the Levels of Selection*, Oxford: Oxford University Press.
- Okasha, S. [2016]: ‘The Relation between Kin and Multi-Level Selection: An Approach Using Causal Graphs’, *British Journal for the Philosophy of Science*, **67**, pp. 435–470.

- Pearl, J. [2009]: *Causality: Models, Reasoning and Inference*, Cambridge: Cambridge University Press, 2nd edition.
- Price, G. R. [1970]: ‘Selection and Covariance’, *Nature*, **227**, pp. 520–1.
- Price, G. R. [1972]: ‘Extension of Covariance Selection Mathematics’, *Annals of Human Genetics*, **35**, pp. 485–490.
- Queller, D. C. [1984]: ‘Kin Selection and Frequency Dependence: A Game-Theoretic Approach’, *Biological Journal of the Linnean Society*, **23**, pp. 133–43.
- Queller, D. C. [1992a]: ‘A General Model for Kin Selection’, *Evolution*, **46**, pp. 376–80.
- Queller, D. C. [1992b]: ‘Quantitative Genetics, Inclusive Fitness and Group Selection’, *American Naturalist*, **139**, pp. 540–58.
- Queller, D. C. and Strassmann, J. E. [2009]: ‘Beyond Society: The Evolution of Organismality’, *Philosophical Transactions of the Royal Society B*, **355**, pp. 1647–1655.
- Rankin, D. J., Mc Ginty, S. E., Nogueira, T., Touchon, M., Taddei, F., Rocha, E. P. C. and Brown, S. P. [2011]: ‘Bacterial Cooperation Controlled by Mobile Genetic Elements: Kin Selection and Infectivity Are Part of the Same Process’, *Heredity*, **107**, pp. 279–281.
- Ridley, M. and Grafen, A. [1981]: ‘Are Green Beard Genes Outlaws?’, *Animal Behaviour*, **29**(944-955).
- Robb, D. and Heil, J. [2014]: ‘Mental Causation’, in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, spring 2014 edition.
<<http://plato.stanford.edu/archives/spr2014/entries/mental-causation/>>
- Sober, E. and Wilson, D. S. [1998]: *Unto Others: The Evolution and Psychology of Unselfish Behaviour*, Cambridge, MA: Harvard University Press.
- Spirtes, P., Glymour, C. and Scheines, R. [2000]: *Causation, Prediction and Search*, Cambridge, MA: MIT Press, 2nd edition.
- Sterelny, K. [2013]: ‘Life in Interesting Times: Cooperation and Collective Action in the Holocene’, in K. Sterelny, R. Joyce, B. Calcott and B. Fraser (eds), *Cooperation and Its Evolution*, Cambridge, MA: MIT Press, pp. 89–108.

- Traulsen, A. [2010]: ‘Mathematics of Kin- and Group-Selection: Formally Equivalent?’, *Evolution*, **64**, pp. 316–323.
- van Veelen, M., Garcia, J., Sabelin, M. W. and Egas, M. [2012]: ‘Group Selection and Inclusive Fitness are Not Equivalent; the Price Equation vs. Models and Statistics’, *Journal of Theoretical Biology*, **299**, pp. 64–80.
- Wenseleers, T., Gardner, A. and Foster, K. R. [2010]: ‘Social Evolution Theory: A Review of Methods and Approaches’, in T. Székely, A. J. Moore and J. Komdeur (eds), *Social Behaviour: Genes, Ecology and Evolution*, Cambridge: Cambridge University Press, pp. 132–58.
- West, S. A., Griffin, A. S. and Gardner, A. [2007]: ‘Evolutionary Explanations for Cooperation’, *Current Biology*, **17**(16), pp. R661–R672.
- Williams, G. C. [1966]: *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*, Princeton, NJ: Princeton University Press.
- Wilson, D. S. [1975]: ‘A Theory of Group Selection’, *Proceedings of the National Academy of Sciences USA*, **72**, pp. 143–146.
- Wilson, D. S., Pollock, G. B. and Dugatkin, L. A. [1992]: ‘Can Altruism Evolve in Purely Viscous Populations?’, *Evolutionary Ecology*, **6**, pp. 331–341.