Biol Philos (2011) 26:501–515
DOI 10.1007/s10539-011-9246-6

DISCUSSION NOTE

Agents and acacias: replies to Dennett, Sterelny, and Queller

Peter Godfrey-Smith

Received: 12 January 2011 / Accepted: 12 January 2011 / Published online: 12 March 2011

Abstract The commentaries by Dennett, Sterelny, and Queller on *Darwinian Populations and Natural Selection* (*DPNS*) are so constructive that they make it possible to extend and improve the book's framework in several ways. My replies will focus on points of disagreement, and I will pick a small number of themes and develop them in detail. The three replies below are mostly self-contained, except that all my comments about genes, discussed by all three critics, are in the reply to Queller. Agential views of evolution, discussed by Queller and Dennett, are addressed in my reply to Dennett.

Keywords Evolution · Genes · Symbiosis · Individuals

Dennett

Dennett locates two main points of disagreement. One is over the role of information in biology, the other over concepts of agency. I'll begin with agents.

DPNS opposes a view of evolution which sees it as a contest between agents pursuing goals. The opposition here is not to selection, adaptation, or optimality, but only to the idea that evolutionary processes themselves are contests between agents. Dennett and Queller both see themselves as defending the agential approach. Some of their arguments are defenses of optimality and adaptation, which are not under attack, but the book was probably not clear enough about what the target is. Once the target has been made clear, I think that Dennett will want to defend it. So clarification does not leave everyone on the same page. Dennett's insistence that a "Cui bono?" question is "central to all evolutionary thinking," for example, is something I am against.

P. Godfrey-Smith (✉)
Department of Philosophy, Harvard University, Cambridge, MA 02138, USA
e-mail: pgs@fas.harvard.edu

The agential view is related to several other controversial forms of description in biology. One is the idea that organisms are akin to products of design, akin to artifacts Lewens (2005). Another is a picture of organisms as adapted to their environments in a lock-and-key manner (Lewontin 1985). A third is the view that organisms will approximate optimal solutions to problems posed by their ways of life (Grafen 2007). These ideas all make subtly different contributions to biology, and probably have different mixes of good and bad effects. The view that is opposed in *DPNS* is the view of evolutionary processes themselves as contests between underlying agents that persist through the material flux. This sort of agential description shades into other members of this family, and into ordinary selectionist thinking.[1]

One possible defense of any of these views is purely pragmatic: "this method has worked so far, so let's keep using it." One cannot give a principled argument against a heuristic understood in this way. If the results are good, that is all that matters. Both Dennett and Queller do appeal to pragmatic arguments, and a response would need a thorough accounting of good and bad outcomes in empirical work guided by the heuristic, with comparison to results gained through alternative methods. *DPNS* does not, as Dennett says, contain such an accounting.[2] But I think the agential perspective is usually seen as closely connected to theoretical considerations, including foundational work. Then more general arguments can be given against it.

To see the distinctive role of agential thinking, we need to look for cases where it takes us beyond other kinds of selectionist thinking and makes its own contribution. Here is an example. *DPNS* discusses the evolution of altruistic and cooperative behavior. One summary idea I accept is that costly forms of altruism can survive if the benefits of altruistic acts fall mostly on those with a tendency to pass the behavior on. Suppose this general idea is embedded within an agential view of evolution. Then the idea becomes something like this: altruistic acts can survive in an evolutionary context because there is an underlying selfish beneficiary which persists and hoards the benefits of these acts. The existence of this beneficiary answers Dennett's "Cui bono?" question.

But altruistic behavior can persist without an underlying selfish beneficiary. An illustration is seen in a recent paper by Fletcher and Doebeli (2009). Assume haploid asexual organisms varying at two loci at which there are alleles causing

---

[1] Dennett says "Ironically, PGS himself finds it irresistible to engage in agential talk when explaining and defending his own best ideas," but I think Dennett casts this net too broadly, and includes adaptationist description and sometimes ordinary causal talk. Here I have inserted comments after his italicized examples. "Here are a few examples among many, in PGS's discussion of de-Darwinizing: "Their independent evolutionary *activities* are c*urtailed, constrained*, or *suppressed* by what is happening at the higher level" (p. 122), [ordinary causal description, with some metaphor, but not agential in the relevant sense] "Another way of dealing with this *problem* [point granted] would be for one member of the collective to *prevent* [ordinary causal description] reproduction altogether by other individuals…" "This can be put more explicitly in terms due to Calcott (2008). For a transition to occur there must somehow be both the *generation of benefit* and the *alignment of reproductive interests*." (p. 124) [point granted, and I do think Calcott's terms are helpful here] So Calcott is praised for putting it "more explicitly" instead of condemned for indulging in such metaphorical agential talk."

[2] Dennett compares the accounting project in Francis (2004), which I think much more highly of than he does.

production of a useful chemical which is made available to others in the neighborhood. The producer alleles are A and B, with a and b as non-producing free-riders. Assume that the AB combination is lethal, as the two pathways that produce the chemical interact destructively. So the population has ab, Ab, and aB individuals. Ab and aB individuals pay a cost for producing the public good. But assume assortment of types, such that Ab and aB interact with each other and ab meets itself. A and B will both proliferate, even though "the help which carriers of A provide never goes to other carriers of A and the help that carriers of B provide never goes to other carriers of B" (Fletcher and Doebeli 2009, p. 16).[3] If a person thought that for altruism to survive, there must be a selfish underlying entity, here they would encounter no selfish beneficiary, and would expect altruism not to survive. But assortment of the two altruist types solves the problem.

A response might be that this is a case of "reciprocal altruism" (Trivers 1971), of the sort also seen in symbioses. This is often regarded as part of the same package of ideas as the agential approach (Dawkins 1976). Here we again reach the problem of working out the distinctive contribution of agential description as opposed to its relatives. I see reciprocity as part of ordinary selectionist thinking, not the special kind which involves the introduction of an underlying agent. A person might also say that both A and B are selfish beneficiaries, or the partnership of A and B is. But to say the former is just to re-express the fact that both A and B do well. To call them a "partnership" is to add nothing that is true to the claim that there is assortment in the population that brings A and B together and each does something that benefits the other.

Turning to the second of Dennett's main themes, Dennett advocates an information-based approach to many issues in biology. Via George Williams, he holds that genes are not "strands of nucleotides" but "the information carried" by those material objects. This affects his view of reproduction, inheritance, and evolutionary processes. In contrast, I do see genes as strands of nucleotides. More exactly, *if* genes are real enough to be Darwinian individuals, then it is strands of nucleotides that are Darwinian individuals—the status of genes as objects will be discussed in my reply to Queller. When one gene is a copy of another, that is a similarity between two separate Darwinian individuals. The same goes with cells, organisms, and reproducing groups. As Darwinian individuals, these are all particular material things. Organisms and cells are not abstract, not made of information, and genes are treated the same way.[4]

Dennett regards reproduction as an informational matter. He gives an argument in which ordinary human reproduction is compared with an artificial case. Suppose we sequence two human genomes, scramble their sequences in a meiosis-like manner, and then construct a physical genome with the results and use it to create a

---

[3] I follow Fletcher and Doebeli's description here though talk of alleles at loci is problematic in the absence of sex.

[4] Sometimes genes *seem* to be more abstract than this because they are not being treated as Darwinian individuals, or other objects, at all. Genetic *properties* are often used to characterize organisms. When this is being done, organisms are being treated as particular objects, perhaps as Darwinian individuals, and genetic features, such as *being AA*, are used to identify a similarity between them. This is not treating genes as Darwinian individuals.

new person. For Dennett, the two people with the genomes we started with are parents of the new individual. I accept this conclusion (adding that the egg provider is a parent too). I also accept some informational or quasi-semantic descriptions of gene action. There is a reasonable sense in which genes do *code for* proteins. But reproduction per se is not an informational matter. Dennett thinks that once one accepts that "'artificial inheritance' can be accomplished by interposing an information-processing device (and nothing else) in the normal process…, that is as good as a proof that the normal process of reproduction is fundamentally an information-transmission process." I reply that the fact that something *can* be achieved with the aid of information-processing machinery is not proof that the task is intrinsically an information-related one. We might use information-processing machinery to do something that in other cases has a different character.

There is a weak sense of "information," used in information theory, in which anything in the world that can take on various different states is a *source* of information, and any other variable whose state is correlated with the first one *carries* information about it. In this sense, clouds carry information about impending rain and footprints carry information about who has walked by. That does not make clouds into informational objects, abstractions as opposed to collections of water droplets, or even into messages. The same applies to genes. Due to the biology of inheritance, parents carry information about features of their offspring. Inheritance is an informational matter in *that* sense, but so is almost any other causal process.

How might we tell whether inheritance is informational in a stronger sense? Dennett, I as understand him, thinks that certain kinds of results can only be achieved by the transmission of some kind of representation. That I take to be part of the moral of the story he tells about spying and submarines.

> We learn that our enemy has somehow *obtained information* about the design of our new submarine. Did a spy copy the blueprints onto pieces of paper and smuggle them across the border, or encode a recipe for the design into patterns of dots and dashes that were sent by radio, or memorize the blueprints and walk across the border, or…..? We might be able to discover that the information moved from Bethesda to Beirut to Bern to Bonn to Baghdad and still not know what physical medium or what code was involved in each leg of the journey.

Dots and dashes are one way to do it, but there may be others. Suppose someone knows that our submarine's design is strongly constrained by a particular technical problem. Many other features will be determined by the way it is solved. Their agent might steal blueprints, but might instead smuggle out a crucial *piece* of the structure, which can be used to infer the whole. ("This looks like a triple-ported Aston-sprocket… but it's been modified… Aha!") In one sense the agent has provided information to our enemy. He has made it possible for them to "re-create structure." But he did not do that using a representation or "code".

Sperber and Wilson, in *Relevance* (1986), their classic in the philosophy of language, argue that work on language has been dominated by the "code model," in which everything that a hearer gets out of a sentence must have been encoded inside it. People tend to assume that this is the only format that is even possible for a theory of language. But this is not the only option. Communication sometimes works by speakers making an inference problem more tractable for hearers. Sperber and Wilson think this is distinct from a code-based explanation (though often working alongside it). In an earlier paper Godfrey-Smith (2000) I imagined a biological inheritance system in which samples of proteins, rather than DNA specifications of proteins, were sent over generations. The offspring runs off new proteins from the samples received from the parents. The device-smuggling case above, the inheritance thought experiment in my earlier paper, and the inference-based linguistic phenomena described by Sperber and Wilson all have in common that they are ways of achieving a result that *can* be achieved by sending encoded messages, but *need* not be achieved that way.

Dennett regards me as a "puritan" about information. History certainly shows how hard it is to avoid being either puritan or libertine.[5] I think that recently it has started to become clear how one might make good naturalistic sense of a large family of semantic phenomena. I have in mind the development of what can be called "sender-receiver" models (Lewis 1969; Millikan 1984; Skyrms 2010). The aim is to give a theory of the ways in which signs operate in the coordination of two kinds of entities, senders and receivers (in broad senses of those terms). Not every sender-receiver interaction involves representations, as the examples above show, so more has to be said. But even when those problems have been solved, this model brings with it a constraint: a sign is only as real as its reader or user. This applies to internal representations as well as to external ones.

A relevant discussion of this constraint can be found in an earlier paper by Dennett, written at a crucial stage in the history of cognitive science.

> The only psychology that could possibly succeed is neo-cognitivist, which requires the postulation of an internal system of representation. However, nothing is intrinsically a representation of anything; something is a representation only *for* or *to* someone; any representation or system of representations requires at least one *user* of the system who is external to the system. Call such a user an exempt agent. Hence, in addition to a system of internal representations, neo-cognitivism requires the postulation of an inner exempt agent or agents—in short, undischarged homunculi. Any psychology with undischarged homunculi is doomed to circularity or infinite regress, hence psychology is impossible. Dennett (1978, p. 101).

Dennett's response to this dilemma is to accept that "perhaps the *prima facie* absurd notion of self-understanding representations is an idea whose time has come, for what are the "data structures" of computer science if not just that: representations that understand themselves?" (p. 102). I think this is the wrong interpretation of the message of computer technology. In standard computers the distinction between representation and reader is present in the hardware of the machine, in the distinction between memory and processor. The message of computer science for

---

[5] There is a sense in which Dennett is both at once, given his interpreter-dependent treatment of semantic properties.

psychology is not the possibility of self-reading representations, but the power of simple read-write devices attached to an organized memory (Gallistel and King 2009).

This framework can also be employed when thinking about genes as carriers of information; this case, too, can be addressed within a sender-receiver model. That is being attempted by Shea (2007) and by Bergstrom and Rosvall (2011). One problem (which I outline in my commentary on Bergstrom and Rosvall 2011) is that identifying a receiver for a genetic message when describing inheritance across generations is not straightforward. If that view can be defended, genes as they operate in inheritance would be information-carrying devices in the stronger sense that fits a sender-receiver model. If not, the language of "information" could still be applied to genes and the relations across generations, but only in the weak sense of this term that can be used in the description of any causal process.

Sterelny

DPNS contrasts "classical" and "replicator" analyses of natural selection. The classical view holds that evolution by natural selection will take place in any population in which variation, heredity, and fitness differences are found (Lewontin 1970). The replicator analysis holds that two roles must be occupied in any evolutionary process. There must be *replicators*, things which are copied in a reliable way, and *interactors* (or *vehicles*), things which interact with the environment in a way that determines the rate at which different replicators are copied (Dawkins 1976; Hull 1980). On the classical view there is no such bifurcation of roles. There is just one kind of entity needed, which I call a "Darwinian individual." These are the things which exhibit variation, heredity, and fitness differences.

I criticized existing versions of both analyses, but then built on the classical view. Sterelny thinks that the replicator/vehicle model is better than I realize. It need not be seen as a complete account, but the framework of *DPNS* allows for a pluralist approach, in which some cases are handled with a refinement of the classical view, and others are handled with the replicator/vehicle view. (I will mostly follow Sterelny's terminology here in talking of "vehicles" rather than "interactors," though I think Hull's concept and terminology are better.)

In *DPNS* the attention was on the replicator side of the Dawkins-Hull framework; once we understand why replicators are not needed, there is no motivation for the bifurcation of roles that introduces a vehicle or interactor concept. Sterelny thinks that I underestimate the importance of both replicators and vehicles. My reply to Queller will discuss replicators. In this reply to Sterelny I look at the vehicle side. Sterelny argues that we need this concept because of the evolutionary role of various entities that are not Darwinian individuals, but do fit the concept of a vehicle well.[6]

Sterelny holds that "extended phenotype" cases are the source of the problem. The cases on his list that I see as especially important are those that involve

---

[6] A view of this kind is also being developed by Austin Booth (forthcoming).

cooperation between members of different reproducing populations. One example is the symbiotic association between ants and bullhorn acacia trees. The trees build hollow structures which house ants that guard the tree. In some cases the trees also make structures that feed the ants. Sterelny claims that the combination, tree-plus-ants, is an adapted unit, built by evolution. But these ant-tree combinations are not Darwinian individuals. They do not stand in parent-offspring relations, except perhaps in a marginal sense. Trees have offspring trees, and ants (and ant colonies) have offspring ants (and colonies), but an ant-tree combination does not reproduce as a unit. It might sometimes be that the ants in a particular tree are the descendants of ants who lived in a parent of that tree, but that would be accidental. The most we can apparently say is that one ant-tree combination, X, is the offspring of another ant-tree combination, Y, if either the ant part or tree part of X (or both) is the offspring of the ant part or tree part of Y. Even this makes the association sound tighter than it often is, as one colony may be supplanted by another and in some cases colonies of more than one species may occupy a tree at the same time (Wheeler 1913; McKey 1984; Palmer et al. 2000; Raine et al. 2004).

Another relevant case here is the squid-vibrio symbiosis, in which colonies of bacteria provide luminescence for bobtail squid (Nyholm and McFall-Ngai 2004). The initial bacterial colony is picked up from the seawater while the squid is young and generally contains a half dozen to a dozen cells.[7] So as with acacias and ants, one squid-vibrio combination is not the parent of another, though squid beget squid and vibrios beget vibrios.

Sterelny holds that "if we recast these evolutionary episodes into the framework of Darwinian Populations, we have to choose between two suboptimal representations." One option tracks the evolution of the clear Darwinian individuals—the ants and the acacias—and does not treat the combination as an entity evolving through natural selection. The other option is to treat the combination as an evolving unit. Then the problem is that we have a complex and adapted structure produced by a marginal Darwinian process. "Novelty and complexity is sometimes not novelty and complexity in a paradigm Darwinian individual. The replicator-vehicle model captures that fact."

I agree that these cases raise special issues and an advocate of the replicator-vehicle view has what looks like a good way of describing them. An ant-acacia combination can be seen as an interactor or vehicle. Vehicles need not be reproducers in their own right. They may arise from the combined actions of members of many replicator lineages, and can serve replicators of several kinds. How might such cases be handled within the Darwinian populations framework?

Focusing on the ant-acacia case, I choose the first option Sterelny offers. Ants are Darwinian individuals, as are ant colonies. So are acacias. Evolution has shaped both sides in such a way that ant-acacia combinations are reliably produced, even though there is not *re*production by ant-acacia combinations. We have an adapted unit that was not produced by a Darwinian process acting on those units.

---

[7] Ruby and Asato (1993) did find that a single bacterium can suffice to establish the symbiosis in a young squid. Wollenberg and Ruby (2009) conclude that the usual number is six to twelve, one or two for each of six "crypts" in the developing light organ.

In *DPNS* I distinguished *origin explanations* and *distribution explanations*. When we give a distribution explanation we assume the existence of a set of variants in a population, and explain why they have the distribution they do. An origin explanation is directed on the fact that a population has come to contain individuals of a particular kind *at all*. In the light of Sterelny's commentary, it is possible to distinguish two kinds of origin explanations that are relevant in an evolutionary context. I will call these *direct* and *co-evolutionary* origin explanations. In the direct case, a single Darwinian population changes over time in such a way that new kinds of individuals arise. In a co-evolutionary case, the entities whose origination we are explaining are not Darwinian individuals at all (at least at first). Rather, two or more Darwinian populations produce Darwinian individuals which interact in some new way that yields the combined or collective entities we are explaining.

This process may involve the evolution of new kinds of individuals within those Darwinian populations, or it might merely involve changes in the distributions of features in those populations, where these changes in distribution lead to novel combinations of entities, drawn from each population, becoming likely to appear.

In Sterelny's ant and acacia case, the unit we encounter is the product of evolution in at least two different Darwinian populations, acacias and ants. (I will be vague from here about whether ants, ant colonies, or both are the relevant Darwinian individuals on the ant side). No direct origin explanation of the kind discussed in *DPNS* is possible. But a co-evolutionary origin explanation can be given. Explanations can be given separately for the ant traits and acacia traits that, when combined in the right way, give rise to the ant-acacia combination. In these explanations, each population operates as part of the environment for the other.

A co-evolving conglomerate at one time can give rise to a Darwinian individual at another. A clear example is the eukaryotic cell, and a less clear example is lichen. This could happen, in principle, with the ants and acacias. Suppose the ants in an ant-acacia complex start to locate and protect the offspring trees of their current mature tree, detecting these offspring trees by chemical means. When the new tree is large enough, an offshoot of the mature ant colony moves in. This would be the beginning of an "egalitarian transition," in Queller's terminology (1997). (This would sound completely far-fetched if ants were not doing equally far-fetched things already, such as herding aphids for food and grooming the aphid eggs in their nests: Matsuura and Yashiro 2006). If the two parties became tightly fused in a way that included mutual reproductive dependence, then when the combination continued to evolve new traits, the origin explanations that would be given would be direct rather than co-evolutionary. We might expect a system like that to produce more adaptive complexity than the current ant-acacia arrangement. Indeed, some ants are not very kind to their acacias: "*C. nigriceps* Emery… has the peculiar behavior at our study site of removing virtually all axillary buds except at swollen thorns, effectively sterilizing the tree" (Palmer et al. 2000).

So I accept that these cases initially seem to fit the replicator/vehicle model well. But the Darwinian populations approach, with additions prompted by Sterelny's objection, can handle them. What is the remaining difference between the analyses? One difference is that the replicator/vehicle analysis adds that an ant-acacia combination is a unit of selection, something that can do well or badly in a

Darwinian context, and in the same sense that applies to sexually reproducing (but non-replicating) organisms like us. As in *DPNS*, I hold that it is a mistake to see things that do not reproduce as units of selection.

Once the distinction between direct and co-evolutionary origin explanations has been made, it starts to make an appearance in many different places. Another topic Sterelny discusses is the role of parameter *S*. This parameter represents the extent to which differences in reproductive output in a population are due to intrinsic differences between individuals rather than extrinsic ones. (The intrinsic properties of an object are those that do not depend on the existence and arrangement of other objects.) I argue that in paradigm Darwinian processes, *S* is high. Reproductive differences are due largely to intrinsic differences. For an example of a reproductive difference that is a low-*S* affair, consider two internally very similar individuals, one who happens to be struck by lightning while the other goes on to have a large family.

Sterelny has many doubts about the role of *S*. He is not even sure that "the concept of an intrinsic trait is well-defined," at least in some contexts where it matters:

> What of those traits whose development is richly dependent on environmental input? The food preferences of a rat, for example, depend on maternal imprinting and early experience, as well as the rat's own genetic endowment. So are preferences intrinsic?

In reply, a preference is an intrinsic trait. It might be a causal consequence of some features that are not intrinsic, and this may affect its heritability, but it is still intrinsic.

Another of Sterelny's examples is more awkward for me and leads to some interesting points. Consider a bird species with an elaborate courtship display that includes behaviors from both sexes. Sterelny speculates that the following might often be true: "Neither sex has a mental template of the whole routine [produced by that sex]. Rather, each bird has a linked set of responses triggered by the previous step of its partner." Suppose we have a case like this. Is the *ability to produce the male courtship display* an intrinsic property of a male bird? We might initially have thought that this was an intrinsic trait, but if the responses made by X's partner are wrong, X cannot perform the display. All X has is the ability to follow a set of conditionals: if I see move 1, I do move 2; if I see 3, I do 4. The other bird follows a different and interlocking set of rules. Provided that one side can get the process started, all the moves will follow in the right order. If it is true that the male has *the ability to produce the male display*, this is not true in virtue of the male's intrinsic properties.

What follows from this? Is it a problem for my view because the male's capacity is a *bona fide* trait that I must treat as questionable? I think, instead, that when we learn these things about the bird's display, it leads us to re-think our description of the trait and its evolution.

The total courtship sequence characteristic of the species, a process involving moves from both partners, is clearly something whose origin explanation is not direct. In that way it is like the ants and acacias. Evolution builds a pair of traits, one expressed in males and one in females. No individual produces both. When the two traits are expressed together, the result is the coordinated courtship display. One might think that at least *one* side of the birds' courtship display, the sequence of behaviors produced by one of the two sexes, is a trait that has a direct origin explanation. But in the case Sterelny is describing, each sex does not have the capacity to produce their side of the display without the prompting of each move by the partner. The temporal ordering of each partner's moves is dependent on the other partner. Then what we think of as "the male's display" does not have a direct origin explanation either; instead it is one side of a co-evolved unit that is not a trait of any Darwinian individual.

Queller

Queller looks closely at genes and the "gene's eye view" of evolution. He agrees that replicators of the sort Dawkins describes are not necessary for evolution by natural selection. For Queller as for me, a simple in-principle argument suffices: if reproduction exhibits some degree of heritability, there can be a response to selection and hence adaptation, whether or not the system contains replicators.

Queller also offers a reconciliation of the "classical" view (based on heritable variation) and the replicator view: "I have never understood why Dawkins could not view individuals [organisms] as loose but good-enough replicators. They make copies of themselves, not especially faithful copies, but faithful enough that selection on parents produces some change in the offspring generation." I don't think this move is easily available to Dawkins. In Dawkins' view of evolution it is central that there are persisting entities that compete with each other. Material objects come and go, but replicators may live on in the form of their copies. Whether or not reproduction is high-fidelity, this picture is much more natural if reproduction is asexual. Asexual reproduction creates a tree in which every individual is at the end of a unique path going backwards in time. When a parent has many offspring asexually, it is possible to regard each offspring as "more" of the stock of the parent. When reproduction is sexual, however, every individual is a meeting of two lineages—actually more than two, because multiple lineages came together in each of its parents. It is not at all straightforward to regard the new individual as a mere continuation, a case of more-of-the-same, with respect to *both* of its parents, if the parents themselves differ. So organisms like us cannot be "loose but good-enough replicators," to use Queller's phrase, unless the link between replication and persistence is broken. Organisms can be *reproducers*, however, because reproduction, in the sense used here, is not a matter of making a continuation of the parent, but making something new, and making something new can be a joint effort.

Genes are a contingent part of the evolutionary machinery, Queller accepts, but undeniably an important part, and the core of his commentary is an application of the framework of *DPNS* to genes. I argued in the book that genes are generally marginal Darwinian individuals. Queller is not convinced. He runs through the tests that I argue should be applied to any candidate Darwinian individual. Genes reproduce, and do so with high heritability (*H*). They vary (*V*). They engage in reproductive competition in the sense measured by α. Is their fitness landscape reasonably smooth (*C*), and are fitness differences due in large part to intrinsic features (*S*)? Yes in both cases. DNA sequences affect organisms' phenotypes, and genes' fitness landscapes should be smooth "to a similar degree as the individuals they are a part of, since the individual fitness landscape derives from its genes."

So far so good, but we have not yet worked out what genes in the relevant sense *are*. If genes are Darwinian individuals, we should know at least roughly where one ends and another begins. We should be able to work out, at least in principle, how many of these individuals are present in a given system. In *DPNS* I asked how we might determine how many individuals are present in a Darwinian population of human genes, and argued that there is no answer. There is a number of human nucleotides, a number of chromosomes, a number of cells, and of organisms, but no number of genes in the evolutionarily relevant sense. When I say there is "no number," I do not mean that no count can be given; the problem is more that too many can. Williams (1966) and Dawkins (1976) say that a gene, for evolutionary purposes, is any stretch of DNA which is subject to positive or negative selection pressure of a degree "several or many times its rate of endogenous change," where endogenous change includes both mutation and crossing-over. Some stretches of DNA will pass a "several" test but not a "many" test, and when the environment changes, shifts in selection pressure will instantaneously wipe out some genes and create others. Many genes will overlap, or exist inside others. A gene that codes for an important protein will contain many sites with slightly different degrees of selective pressure bearing on them.

Queller offers a response to this problem. The way to recognize genes as Darwinian individuals is to treat them, in general, as single nucleotides.

> The word "gene" has notoriously been used for many purposes. I think the meaning that matters most here a population geneticist's gene: a site that varies, especially if it varies in a way that makes a phenotypic difference that selection could work on. Let's make it the minimal such site, which will often be a single nucleotide difference, unless it is a larger deletion or insertion.

This resolves ambiguities over counting. Queller's genes have definite boundaries and no overlaps. His next move is to hold that separate Darwinian populations exist for each locus, where, again, a locus is the site occupied by a single nucleotide.

I will spend some time looking at the consequences of this view. I will set aside the role of insertions and deletions, and look at the possibility of viewing nucleotides as Darwinian individuals.

The first thing to note is that there are only four character states possible for members of each population: C, A, T, and G.[8] There are no traits of Darwinian individuals that are combinations of a number of variable characters. Mutation can replace a Darwinian individual of one kind with another, and the frequencies of the four types can go up and down, but no other kind of change is possible in such a population. This is Darwinian atomism.

---

[8] Or five if methylated C is a different option from unmethylated C.

DPNS distinguished between *origin explanations* and *distribution explanations*, where the former explain the appearance of new variants in a population and the latter explain why some variants are common and some are rare. In a genetic Darwinian population of Queller's kind, no origin explanations can ever be given, (except in the earliest stages). Once all four nucleotides are present, no new traits ever arise. Larger pieces of genetic material—such as cistron-sized pieces—can evolve in a gradual way, with each substitution building on its predecessors and new combinations arising from old. But those larger entities are made up of large numbers of Darwinian individuals, each located in a different population. In my reply to Sterelny, I added to the framework of *DPNS* by distinguishing *direct* from *co-evolutionary* origin explanations. The aim was to handle cases in which novel structures arise through evolution in several different Darwinian populations. Ant-acacia symbioses are examples. In Queller's framework, almost all origin explanations are co-evolutionary; even the ability to code for a single protein requires the coordination of many Darwinian individuals.[9]

This does not mean that direct origin explanations for traits like eyes and immune systems cannot be given at all. They can be given by treating organisms as Darwinian individuals. Queller does not argue for an *exclusively* gene's eye view of evolution. Population genetics is very much a science of distribution explanations.

I do not see the points above as arguments against Queller's view yet, especially given that the deficiencies in a gene-level description can be made up by moving to organisms. The next issues I will raise are more of a problem. At an initial stage in Queller's discussion he said that genetic Darwinian individuals will often have reasonably high scores for *S* and *C*. But his decision to identify genes in most cases with single nucleotides changes this. Queller notes this fact himself for the case of *S*. When one genetic individual reproduces and another does not, this may be partly due to an intrinsic difference between them—being A rather than T—but a huge number of realized fitness differences will be due to extrinsic factors. Setting aside "environmental" factors in the usual sense, whenever one individual does well and one badly because of what is going on at other loci (including the nucleotides next door), realized fitness is affected by extrinsic matters.

We also need to look again at *C* (for "continuity," which I italicize here to distinguish it from the C nucleotide). *C* is a measure of whether small changes to individual character lead to small changes in fitness. Queller said initially that genes will often have respectable *C* scores, but this changes once genes are the size of nucleotides. Then it is hard to even ask the question about continuity, because a genetic population has just four kinds of individuals, with no intermediates between the different states. There is one continuity-related feature that might be recognized. The purines (A and G) could be seen as more similar to each other than each is to the pyrimidines (C and T). Pyrimidines would also be internally similar. Whether this has consequences for fitness depends on the structure of the genetic code and on what various proteins do. Suppose transitions (moves between purines, or between

---

[9] I am not sure how Queller intends his analysis to apply to asexual organisms, such as bacteria, where the entire genome is the usual unit of replication. Here I assume there are not separate Darwinian populations for each nucleotide-sized locus, as facts about competition at a locus are dependent on meiosis.

pyrimidines) generally have fewer phenotypic consequences and hence fewer fitness consequences than transversions (movements between a purine and a pyrimidine). That would make for some degree of continuity. But that seems to be just about the only question that can be even asked about *C* for this Darwinian population. This is because of the atomism of Queller's version of the gene's eye view.

I will address two other topics in this area. One is the only issue on which there is some misunderstanding. Genes, it is clear, are scaffolded reproducers. Their reproduction depends on much machinery that is external to them. Queller attributes to me the view that cases where reproduction is scaffolded are "less Darwinian," and he disagrees with this. That claim about scaffolded reproduction is not part of my view, however. HIV is a scaffolded reproducer, for example, and it is a paradigm evolver (p. 152). Scaffolded reproducers are special because their reproductive capacities have a kind of fragility. But that does not make their evolutionary processes marginal when the scaffolding they require is present.

Finally, Queller notes that I accept a Darwinian description of gene action for some special cases: selfish genetic elements.

> Yet, they partake of all of the supposed non-Darwinian sins of other genes… They fall into about the same region in the state space of Darwinian populations as other genes. What really separates them from other genes is not anything on the list of important Darwinian properties, but the additional simple fact that they *cannot* be folded into a strictly organismal view. To me, that constitutes grounds for *accepting* other genes as Darwinian individuals as well.

This is indeed a problem. If these special cases are properly described in terms of selection on genes then they should be different *in their own right* from the ones where I resist the gene's eye view in favor of an organism-level description. The question should not be which description works "best," so that some cases receive a gene-level description because rival descriptions are especially bad. The question should just be how well genes fit the Darwinian criteria in each case.

Queller's objection might be met in at least some cases. Meiotic "driving" is often the capacity of a large complex on a chromosome that is protected from crossing-over by an inversion (Burt and Trivers 2006). That gives them more integrity than other stretches of genetic material. LINE transposons, another example used in the book, also have clearer boundaries than most genetic elements, because they travel as units. If something of this kind is not applicable in each case, then we have a situation where the cases that seem to *need* a Darwinian description at gene level do not "partake of the same sins" as the cases where I resist or deflate that description.

Queller's exploration of the gene's eye view is illuminating. A trade-off emerges. One option is to resolve problems of counting and boundaries by regarding the genetic Darwinian individuals as the minimal units, nucleotides, which are indivisible for these purposes. Then each population has only four kinds of individuals, problems arise with *S* and *C*, and selection within a single Darwinian population never explains origins. Another option is to return to a view in which evolutionary genes are larger units, perhaps roughly the size of a cistron or promoter. Then the conclusions regarding *S* and *C* that seemed appealing at the start of Queller's discussion can be restored, and direct origin explanations can be given for the shaping of genes with special causal capacities. This view must then confront the problems that genes of this kind are very dubious entities. Defenders of this version of the gene's eye view have often talked about genes as if they were objects that could be counted at least in principle (allowing some vagueness), and which could be treated as comprising entities in a population. *DPNS* argued that this is an illusion. In the light of Queller's discussion, I think that this second option is nonetheless the more promising direction for the defender of the gene's eye view to pursue. This is partly because of the special cases—meiotic drivers, transposons, and other selfish genetic elements. Insofar as they force us to recognize a version of the gene's eye view, it is a view of this second kind rather than Queller's atomism.

References

Bergstrom CT, Rosvall M (2011) The transmission sense of information. Biol Philos. doi: 10.1007/s10539-010-9233-3

Booth A (forthcoming) Selection, symbiosis, and individuality

Burt A, Trivers R (2006) Genes in conflict: the biology of selfish genetic elements. Harvard University Press, Cambridge

Calcott B (2008) The other cooperation problem: generating benefit. Biol Philos 23:179–203

Dawkins R (1976) The selfish gene. Oxford University Press, Oxford

Dennett DC (1978) A cure for the common code. In: Brainstorms: philosophical essays on mind and psychology. Bradford Books, Cambridge, pp 90–108

Fletcher JA, Doebeli M (2009) A simple and general explanation for the evolution of altruism. Proc Royal Soc B 276(2009):13–19

Francis R (2004) Why men Won't ask for directions: the seductions of sociobiology. Princeton University Press, Princeton

Gallistel CR, King A (2009) Memory and the computational brain: why cognitive science will transform neuroscience. Blackwell/Wiley, New York

Godfrey-Smith P (2000) On the theoretical role of 'Genetic Coding'. Philos Sci 67:26–44

Godfrey-Smith P (2010) Senders, receivers, and genetic information: comments on Bergstrom and Rosvall. Biol Philos. doi:10.1007/s10539-010-9206-6

Grafen A (2007) The formal Darwinism project: a mid-term report. Evolut Biol 20:1243–1254

Hull DL (1980) Individuality and selection. Annu Rev Ecol Syst 11:311–332

Lewens T (2005) Organisms and artifacts: design in nature and elsewhere. MIT Press, Cambridge

Lewis DK (1969) Convention: a philosophical study. Harvard University Press, Cambridge

Lewontin RC (1970) The units of selection. Annu Rev Ecol Syst 1:1–18

Lewontin RC (1985) Adaptation. In: Levins R, Lewontin RC (eds) The dialectical biologist. Harvard University Press, Cambridge, pp 65–84

Matsuura K, Toshihisa Y (2006) Aphid egg protection by ants: a novel aspect of the mutualism between the tree-feeding Aphid *Stomaphis hirukawai* and its attendant ant *Lasius productus*. Naturwissenschaften 93:506–510

McKey D (1984) Interaction of the ant-plant *Leonardoxa africana* (Caesalpiniaceae) with its obligate inhabitants in a rainforest in cameroon. Biotropica 16:81–99

Millikan R (1984) Language, thought, and other biological categories. MIT Press, Cambridge

Nyholm S, McFall-Ngai MJ (2004) The winnowing: establishing the squid-vibrio symbiosis. Nat Rev Microbiol 2:632–642

Palmer T, Young T, Stanton M, Wenk E (2000) Short-term dynamics of an acacia ant community in Laikipia, Kenya. Oecologia 123:425–435

Queller DC (1997) Cooperators since life began (Review of *the major transitions in evolution*, by J. Maynard Smith and E. Szathmary). Q Rev Biol 72:184–188

Raine N, Gammans N, MacFadyen I, Scrivner G, Stone G (2004) Guards and thieves: antagonistic interactions between two ant species coexisting on the same ant-plant. Ecol Entomol 29:345–352

Ruby E, Asato L (1993) Growth and flagellation of *Vibrio fischeri* during initiation of the sepiolid squid light organ symbiosis. Arch Microbiol 159:160–167

Shea N (2007) Representation in the genome, and other inheritance systems. Biol Philos 22:313–331

Skyrms B (2010) Signals: evolution, learning, and communication. Oxford University Press, Oxford

Sperber D, Wilson D (1986) Relevance: communication and cognition. Blackwell, Oxford

Trivers RL (1971) The evolution of reciprocal altruism. Q Rev Biol 46:35–57

Wheeler WM (1913) Observations on the central American acacia ants. In: Jordan K, Eltringham H (eds) International conference of entomology, oxford, August 1912, vol 2 Transactions. Hazell Watson and Viney, London, pp 109–139

Williams GC (1966) Adaptation and natural selection: a critique of some current evolutionary thought. University of California Press, Berkeley

Wollenberg M, Ruby EG (2009) Population structure of *Vibrio fischeri* within the light organs of *Euprymna scolopes* squid from two oahu (Hawaii) populations. App Environ Microbiol 75:193–202