

Genetic drift as a directional factor: biasing effects and a priori predictions

Ariel Jonathan Roffé¹

Received: 17 October 2016 / Accepted: 24 May 2017 / Published online: 7 June 2017
© Springer Science+Business Media B.V. 2017

Abstract The adequacy of Elliott Sober’s analogy between classical mechanics and evolutionary theory—according to which both theories explain via a zero-force law and a set of forces that alter the zero-force state—has been criticized from various points of view. I focus here on McShea and Brandon’s claim that drift shouldn’t be considered a force because it is not directional. I argue that there are a number of different theses that could be meant by this, and show that one of those theses—the idea that drift cannot bias populations to be taken somewhere in the evolutionary space from one generation to the next—is actually false. Not only has this thesis been implicitly assumed in the discussion of the force analogy thus far, but it is also commonly found in a wider range of philosophical and biological texts. I argue that correcting this view, and the usual images associated with it, will thereby bring heuristic benefits that impact the force analogy discussion, but that also go beyond it.

Keywords Drift · Directionality · Evolutionary space · Evolutionary forces

Introduction

In the first chapter of his 1984 book, *The Nature of Selection*, Elliott Sober articulated an analogy between evolutionary theory and classical mechanics, according to which the former can also be understood as a theory of forces. I write

✉ Ariel Jonathan Roffé
ariroffe@hotmail.com

¹ CONICET, CEFHIC-UNQ (Center of Studies in the Philosophy and History of Science, National University of Quilmes), UNTREF (Universidad Tres de Febrero), Buenos Aires, Argentina

“articulated” because the idea was hardly new at the time. The analogy was already present (either explicitly or implicitly) in evolutionary theory textbooks and publications, as the use of force language to describe evolutionary phenomena was (and is) common in them. For example, in their textbook from 1970, Crow and Kimura claimed that “[n]atural selection [as mathematically treated by population genetics], like classical mechanics, has both static and dynamic aspects. The statics of evolution [... involve] the relatively stable situation that results from the balance of various opposing forces—mutation, selection, migration, and random fluctuations” (Crow and Kimura 1970, p. 175).¹ What Sober did was to develop this idea in a more detailed manner. Thus, according to him, both theories are analogous in the sense that they construct their *explanans* via a zero-force law—the law of inertia and the Hardy–Weinberg law, respectively—and a set of forces that can alter or bias the equilibrium state characterized by it—selection, mutation, migration and drift, in the evolutionary case—, and which can add up to produce complex effects.²

I take it that the force analogy is interesting for at least two reasons, a metatheoretical and a pragmatic one. Metatheoretically, it could be useful to academics who study science, because (if adequate) it allows us to recognize common explanatory strategies between disparate fields of science (such as physics and biology), which is relevant to an understanding of how science works, and how scientific theories explain. Or, to put it in Earnshaw’s (2015) terms, to recognize places in which the same “modeling strategy” is being used. Pragmatically, it could also interest researchers themselves (biologists and philosophers alike), because it would give us an interesting frame or heuristic to think about evolutionary phenomena, and also to educators of science, if it is (or is not) conducive to better teaching practices of evolutionary theory. This last point is especially important since, as noted above, the analogy is widely used in materials used to teach evolutionary theory. For all these reasons, assessing the adequacy of the force analogy is an important task for philosophers of science.

Within philosophical literature, the analogy has received a number of criticisms that could be interpreted as showing that important disanalogies between these two theories render it (or at least the way Sober presented it) inadequate. One of those objections, which is the focus of this paper, is that one of the evolutionary factors initially identified with a force—genetic drift—should not be considered as such. The reason being that, unlike Newtonian forces and the rest of the evolutionary

¹ The same uses of language can be found in more recent publications, as can be seen, for example, in Gillespie’s textbook, where he states that “[t]he Hardy–Weinberg law describes the equilibrium state of a single locus in a randomly mating diploid population that is free of other *evolutionary forces*, such as mutation, migration, and genetic drift.” (Gillespie 1998, p. 11, *emph. added*) or in Hartl and Clark’s, as they claim that “[t]o deduce the [Hardy–Weinberg] genotype frequencies under random mating, additional assumptions are needed. First, the allele frequencies should not change from one generation to the next because of systematic *evolutionary forces*, the most important of which are mutation, migration, and natural selection” (Hartl and Clark 2007, p. 48, *emph. added*). For more examples, see Pence (2016).

² There are other ways in which an analogy between these two theories has been made. Darwin himself made some remarks in that direction (Darwin 1872, p. 63, 421). For a more recent example (and curiously, partly in response to Sober 2011), Díez and Lorenzano (2013) argue that they are analogous in the aprioricity status of their laws. Other possible ways of making the analogy can be seen in Williams (1980). I will focus here on Sober’s (1984) presentation.

factors, drift is not a *directional* factor (Brandon 2006; McShea and Brandon 2010). That is, for populations of organisms evolving under genetic drift, it is not possible to a priori predict the “direction” (in a sense to be explicated) that allele or genotype frequencies will take.

The goal of this paper is to argue that the last sentence is not equivalent to the one before it; i.e., that we should distinguish between a *directional factor*—one that, in Sober’s terms, has “biasing effects” on allele and genotype frequencies—, and a factor for which we can (reliably) predict a particular direction in an a priori manner.³ Even conceding that the direction in which drift takes a population cannot be predicted a priori (at least not from one generation to the next), I will claim that drift should still be considered a directional factor, since (except in some very particular cases) it can “bias” populations to go one way or the other, even from one generation to the next.

To make this point I will show that a particular claim often made explicitly or assumed implicitly, in both the philosophical and biological literature, is not true: that for a population evolving under drift, the probability that an allele frequency increases is the same as the probability that it decreases in the immediate next generation. In contrast (again, with some qualifications, see below), I will show that for a population whose starting allelic frequency is greater than 0.5 for one of the alleles, the probability that this frequency goes up is higher than the probability that it goes down (and vice versa); I will argue that this must count as a kind of “bias” in Sober’s original sense.

However, there is also some truth in McShea and Brandon’s claims. It is true that the biases mentioned are never large enough to generate reliable predictions in the short term—i.e. even if the probability that an allele frequency increases is greater than the probability that it decreases, none of those “directions” can be predicted with a probability greater than 0.5. Thus, in the short term at least (from one generation to the next), it is actually true that drift’s direction cannot be predicted reliably. Therefore, I will maintain that McShea and Brandon’s objection, understood in a certain way, is potentially *misleading*; claiming (either as a way to object to the use of the force analogy or independently of this matter) that drift is not a directional factor can lead (and has led) to a bad conception of genetic drift, and to the committing of mistakes. Correcting this point will thus have benefits, both metatheoretical and heuristical, as it will give us a clearer picture of the similarities and dissimilarities in the explanatory structures, and avoid those kind of mistakes and inadequate images of drift in the future (which, as I will show, appear even in well-known textbooks). All of this does not necessarily mean that I think drift *should* be equated with a force. Thinking of the analogy as a heuristic, the criterion for being a force becomes an ultimately pragmatic matter, one which I will not

³ Of course I do not mean here an evolutionary factor for which one cannot make directional probabilistic predictions at all. For instance, utilizing drift-only models, one can obviously predict that, for some particular directional event e (say, the event that an allele’s frequency will increase in the next generation), $P(e) = x$. The point is that, with drift-only models, and considering a specific event-type E (short-term directional predictions, see below), then $P(e)$ is never high enough to justify the belief that e will actually take place.

address here in full generality. However, I believe this discussion will allow the pragmatic discussion to move forward in a more fruitful and accurate manner.

For the remainder of this paper, I proceed as follows: The next section introduces in more detail Sober's analogy and McShea and Brandon's objection; it then explicates a number of key concepts needed to make them both more precise. Section 3 distinguishes a variety of possible theses that McShea and Brandon could be proposing. Section 4 focuses on one of those theses—the idea that drift cannot bias populations to be taken somewhere in the evolutionary space—, shows it to be false, and draws some consequences from this fact. Finally, some general conclusions are drawn, and the subject of the adequacy of the force analogy is revisited. A Mathematical appendix adds some more technical considerations regarding the issues explored in the paper.

Explicating some key terms in the force analogy: the particle model

In “Evolutionary Theory as a Theory of Forces”, the first chapter of Sober's 1984 book *The Nature of Selection*, he details an analogy between this theory and classical newtonian mechanics. More specifically, the analogy is between *population genetics* (which is given a special place within evolutionary theory⁴) and newtonian mechanics, as will become clear below. According to him, both theories would be analogous in a sense that could be called “structural”,⁵ which has to do with the way they construe their explanations. That is, both would explain their respective phenomena parting from a zero-force law, which describes the state that an entity will be in, in the absence of forces, and a series of forces which can alter those states, and which can add up to produce complex effects.

In the biological case, Sober claims that forces can be recognized by their “biasing effects” on allele, genotype or phenotype frequencies:

One of the main goals of what follows is to distinguish between different possible causes of evolution [i.e. evolutionary forces]. *All* possible causes of evolution may be characterized in terms of their “biasing effects”. (...) All this is merely to locate evolutionary theory in familiar territory: it is a theory of forces (...) This idea will carry us some distance toward understanding the structure of evolutionary theory—how its models are put together and how

⁴ That is not to say that evolutionary theory consists solely of population genetics. I would argue that this is not true even for microevolutionary theory (see Roffé and Ginnobili, in press).

⁵ This is not the only noteworthy aspect of Sober's analogy. Recently, two authors have claimed that “[t]here are (at least) two different ways in which [classical mechanics and evolutionary theory] may be analogous: (1) Evolutionary forces are like Newtonian forces in the way that they are used to construct mathematical models of the evolution of a system in time. (2) Evolutionary forces are like Newtonian forces in being causes of the temporal evolution of the system” (Hitchcock and Velasco 2014). Both of these aspects appear clearly in the quote below. My treatment here will be more akin to the first perspective, as I will study genetic drift, its directionality, and its similarity with the other factors only as they are modeled by theories like *population genetics*. In other words, my concerns will be mainly metatheoretical/epistemological, not metaphysical.

they make contact with the reality of the living world. (Sober 1984, p. 31, *emph. in original*)

Evolutionary theory's forces would then consist in selection, mutation, migration and genetic drift.⁶ The zero-force law would be the statement that describes what obtains when all these forces are absent, which is, according to Sober, the Hardy–Weinberg law.

A challenge to the force analogy, at least as presented above, was posed by Brandon (2006) and McShea and Brandon (2010), who claim that one of the factors that Sober had identified with an evolutionary force—genetic drift—has a different structural role in evolutionary theory than the rest of the factors, and is therefore not analogous to them (nor to Newtonian forces). One of the arguments they present is based on the idea that drift—again, unlike the rest of the factors—“has no direction”. Since Newtonian forces are vectorial quantities that have both a magnitude and a direction, then drift shouldn't be considered a force (Brandon 2006, pp. 324–325; McShea and Brandon 2010, p. 101).⁷

It should be noted, firstly, that if drift can “move” populations from the state they are in, or from the state that they were expected to be in, then it *must* be moving them in some direction (however one chooses to understand “direction” in this context, I will make this more precise later on). What Brandon and McShea are claiming is not so much that drift doesn't have a direction, but that its direction is not specifiable or predictable a priori; thus, what they claim is that “the direction in which drift takes a population can be determined only after the fact” (McShea and Brandon 2010, p. 101).

They are not entirely clear about what this means exactly, nor about how they reached this conclusion. A necessary first step towards clarifying their objection and assessing its adequacy, is to explicate a number of key concepts present in the force analogy, such as “evolutionary space” and “direction”, that have, so far, been used somewhat intuitively. To do this, McShea and Brandon's own “particle model” (McShea and Brandon 2010, pp. 14–15) can be used.⁸

The model consists of a particle that moves, randomly (with probability 0.5), one vertical step (column), either to the right or to the left, at each horizontal step (row) s (it cannot remain in the same place; from here on, the word “step” is used only to refer to horizontal steps). The probability distribution for the particle movement after four steps have passed is shown in Table 1.

What the particle model does is graphically represent a *Bernoulli sampling process*. This kind of process is one in which a sample of size n is taken, with replacement, from a population (in the probabilistic sense of the word, a set of

⁶ Perhaps one should include some other ones, such as meiotic drive or genetic draft (Skipper 2006); I only mention these four for simplicity.

⁷ The point had already been made by Matthen and Ariew (2002, p. 61). Still, it was McShea and Brandon who fully fleshed out the consequences of this thesis; so, my focus is on their presentation.

⁸ The particle model was originally introduced to show that a non-directional process at one level can have directional effects at the next level (e.g. an increase of the variance at the level of ensembles of populations)—see McShea and Brandon (2010), chapter 2. Nevertheless, I believe it can be fruitfully used to illustrate the authors' conception about drift (the particles themselves can be thought of, according to them, as drifting populations).

Table 1 Probability distribution for McShea and Brandon’s particle model

s_4	0.0625	0.25	0.375	0.25	0.0625
s_3	0.125		0.375	0.375	0.125
s_2		0.25	0.5	0.25	
s_1			0.5	0.5	
			1		
	0	0.25	0.5	0.75	1

objects), which is composed of two types of things. For instance, one can think of the population as a set of marbles of two colors (say red and blue) inside an urn. To put a more concrete example, let us say that there are infinitely many balls in the urn; that in the urn, $p = 0.5$ (with p representing the frequency of the red balls), that is, half of the balls are red; that the sample size is $n = 4$; and that the 4 balls in the sample are taken from the urn simultaneously, not one by one. We may think of this as four different people each grabbing a different ball at the same time.

The idea is that if one wishes to calculate the probability that the sample contains exactly k red balls one can proceed *as if* the balls were being taken successively. For example, for $k = 4$, one can reason as follows: the probability that the “first” ball is red is 0.5. Of the 50% of cases in which the first ball is red, there will be another 50% of those in which the “second” ball is also red. That is, the probability that the “first two” balls are both red is $0.5 \times 0.5 = 0.25$ (notice that this is the number in the right at the s_2 row, in Table 1). In total, the probability that the sample contains all four red balls equals $0.5^4 = 0.0625$ (again, see the rightmost entry in the s_4 row, in Table 1).

Note that the calculation for k other than 0 or 4 is a bit more complex. For example there are four possible ways in which the sample can contain $k = 1$ red balls. Remember that four different persons are each grabbing a ball from the urn at the same time, so each of them can get the red one. Represented as if the sample was taken sequentially, this amounts to summing the probabilities of the samples $\langle R,B,B,B \rangle$, $\langle B,R,B,B \rangle$, $\langle B,B,R,B \rangle$, and $\langle B,B,B,R \rangle$, each of which has a probability of $0.5^1 \times 0.5^3$. The result is 0.25 (see the second column of the s_4 row). The different ways to get a sample with the same composition (i.e. frequency for both kinds of objects) are represented in the particle model as *different paths* leading to the same endpoint in s_4 . More generally, the probability that the sample contains exactly k red balls and $n - k$ blue ones is given by:

$$\binom{n}{k} p^k q^{n-k}$$

An equation known as a *binomial distribution*. The particle model from above gives, at the top row, the same values as a binomial distribution.

Considered in this way, as a graphical representation of a Bernoulli sampling process, one can see that the various parameters in the particle model represent the following sampling parameters. The probabilities of the particle of going right and

left correspond to the frequencies of the two types of objects in the population (which, in turn, correspond to their probabilities of being sampled individually at each “step”). Each step represents the adding of an individual to the sample being taken. Therefore, the number of steps represents the sample size n (in the above case, $n = 4$). This is important because, as was said, the balls in the example are being taken simultaneously, not sequentially, so the different positions of the particle at the intermediate steps ($s_1 - s_3$) do not represent any real (or even possible) entity. They are only there to help us make the calculations (i.e. as was said, we are proceeding *as if* the sample was taken sequentially). In other words, the “steps” s_i represent *logical* steps (in the calculation) rather than *temporal* steps. Only the different possible positions of the particle at s_4 represent the different possible kinds of sample, and the axis values tell us the kind of sample (the frequency of one kind of object) for a particle at s_4 . The axis values do not represent anything at the intermediate locations ($s_1 - s_3$).

The reason this model is relevant for thinking about biological populations of organisms evolving under drift, is that population genetic theory represents generational transitions in terms of sampling processes (for example, Beatty 1984, pp. 188–189 calls them “parental sampling” and “gamete sampling”); and, particularly, drift-only models typically employ binomial distributions (for instance, the Wright–Fisher model utilizes them iteratively). Hence, the particle model can be seen as a representation of the probabilities of the outcomes of population genetic processes in drift-only models.

More precisely, the population may be thought as a population of parental alleles or gametes,⁹ the probabilities of going right and left being the frequencies of an A and an a allele, and the number of steps s_i being the size of the daughter (allele) population at birth (i.e. twice the organismic daughter population size). Again, notice that advancing up a step (for instance, going from s_2 to s_3) does not represent the passing of a biological generation, only at s_4 we have completed a generational transition. Thus, the particle model of Table 1, applied to a biological population, should only be interpreted to indicate what Fig. 1 shows.

The spatial nature of the particle model is very useful for explicating some terms commonly used in the force analogy. For instance, the “evolutionary space” can now be seen as the one-dimensional space represented in the axis of the model; the different positions of this space, therefore, represent the frequency of one of the alleles (let us say, for convenience, the A allele). Every possible two-allele biological population (represented as a particle) occupies a position in this space.¹⁰ A “direction”, which is more naturally understood in a one-dimensional space as “to the right” or “to the left”, can then be understood as “towards the increase of the frequency of A ” or

⁹ A gamete pool is, in most cases, an abstract entity or an idealization, not a spatiotemporally existing one. However, this does not matter. All that is required is that population genetic models do represent generational transitions in terms of those kinds of (statistical, if not biological) populations. That will be enough to make my point that, *according to the usual models*, evolution by drift is a biased process (see below).

¹⁰ Or, more accurately, every possible population of size 4, since the particle model in Table 1. includes a discrete axis. However, one may easily extend this image, and think of a space which has a continuous axis, as in Fig. 1.




Fig. 1 Biological interpretation of Table 1

“towards the decrease of the frequency of *A*”, and the magnitude of drift would tell us how much or how little it increases or decreases. So, for instance, if the population represented by the particle in Table 1 ends up having a frequency of 0.75 at the start of generation $i + 1$, we would say that drift moved it “to the right” and with a magnitude of 0.25 (this is not standard talk in the biological literature, but it may be the best way to understand Sober’s analogy in a precise manner).

McShea and Brandon’s objection: three possible interpretations

In this section, McShea and Brandon’s claim that drift does not have an a priori specifiable or predictable direction is reexamined and clarified, with the aid of the ideas developed in the previous section. I will hold that the claim above could be interpreted as three different theses; in later sections I argue that one of those theses is incorrect.

One way to begin, is to consider the way in which other evolutionary factors may be said to have an a priori predictable direction. Consider, thus, a population evolving by natural selection—let’s say, for simplicity’s sake, one which mates randomly and has discrete, non-overlapping generations—and consider a single *locus* for which there are two allele types (call them *A* and *a*). Let also w_1 , w_2 and w_3 be the fitness coefficients (i.e. probabilities of making it into adulthood) of genotypes *AA*, *Aa* and *aa*, respectively. If, at birth, frequencies for these genotypes are p^2 , $2pq$ and q^2 (p and q being the allele frequencies of *A* and *a*), then, all else absent (including drift), population genetic theory implies that the adult frequencies p^* and q^* for alleles *A* and *a* will be:

$$p^* = \frac{p^2w_1 + pqw_2}{p^2w_1 + 2pqw_2 + q^2w_3}, \text{ and } q^* = 1 - p^*,$$

and genotype frequencies at birth in the next generation will equal p^{*2} , $2p^*q^*$ and q^{*2} . In other words, given the state of a population in a generation i , selection-only models return a precise number that tells us, in a deterministic way, what the state of the population (i.e. allele and genotype frequencies) will be in generation $i + 1$.

Therefore, both the direction and magnitude (in the sense explicated in the last section) in which selection takes a population, can be predicted precisely, if one is able to determine the parameters in those equations. McShea and Brandon's point seems to be that drift-only models cannot do this. But, again, this could mean a number of different things.

First, it could be seen as an argument in favor of the claim that drift models, unlike selection models (and newtonian mechanics' models), can only make *probabilistic* rather than deterministic directional predictions. That is, that one cannot tell *for certain* what the state of the population will be when all else is absent, and therefore drift shouldn't be equated with a force because forces are supposed to be deterministic. This doesn't seem like a particularly interesting claim, since everybody has known from the beginning that population genetics is a probabilistic theory, so nobody expected the theories in question to be analogous in this sense.¹¹

Second, it could be an illustration of a stronger and more interesting claim, that drift-only models do not even allow us to make *good/reliable probabilistic* predictions regarding the directions in which populations will go. An example might help clarify what this second version means. If one tosses a fair coin, then there is a 0.5 probability of it landing heads; that is, one can predict that it will land heads, but knowing that the prediction will come true only about half the time—which doesn't seem like a very good/reliable prediction.¹² In contrast, if the coin is heavily biased (say, the probability of heads is 0.99), then the prediction that it will land heads seems more solid, since it will come true approximately 99 out of 100 times.¹³ I believe McShea and Brandon want to make the stronger point that, with drift-only models, no *good/reliable* predictions (not even probabilistic) regarding direction can be made. This can be seen in their response to an objection to their claim (which can be found, for example, in Barrett et al. 2012, pp. 730–731). Standard population genetics theory implies that, in the long term, drift tends to fix the most frequent allele more often than the reverse. Indeed, if the starting frequency of *A* is 0.99, then there is a 99% chance that *A* becomes fixed. The prediction that *A* will become fixed seems then like a reliable one in this case. More generally, it could be said that drift's direction is something like “towards the fixation of the most frequent allele”.¹⁴

¹¹ Although Pence (2016) has argued that classical mechanics does contain stochastic forces as well, like brownian motion, and that the theories are analogous even in this sense. I will not go into this discussion here.

¹² Of course, with a fair coin, one could also predict that, in a long series of throws, it will land heads approximately half the time and tails about half the time. But that would not be predicting a direction for the coin toss, but the probabilities that the coin goes in different directions.

¹³ The question of the conditions under which a probabilistic prediction is reliable is an important, and very difficult, one. It seems that they depend, at least partly, on the pragmatic context. For example, some have argued that the threshold of reliability should be set as a result of a balance of the risk/cost of getting a false positive vs. the risk/cost of getting a false negative (Rudner 1953). In the context at hand, it seems that being better than a coin toss is at least a necessary condition for reliability. Therefore, I will take having a probability greater than 0.5 as a necessary condition.

¹⁴ A similar point had been made by Stephens, who claims that drift can be thought of as having a direction towards the reduction of homozygosity (Stephens 2004). Brandon's response to this is to claim that, while it is true that drift tends to fix alleles, this is not a directional prediction since it does not say which of the alleles will end up fixing (Brandon 2006, p. 325). For the thesis presented above, this response is not available, since the fixation of a specific allele is being predicted.

If the point McShea and Brandon were trying to make was the weaker one—that the directional predictions one gets from drift are probabilistic rather than deterministic—then their response should be exactly that. The prediction that the most frequent allele will be fixed is only true with some probability, possibly very high, but different than 1. However, their actual response is something different. According to them, it is not drift alone which makes alleles fixed, but drift *plus* absorbing barriers. In their own words:

One of the standard predictions of the theory of genetic drift is that it eliminates genetic variation from populations. The dynamic of this is easy to understand. In figure 6.2 [the particle model applied to a biological population] we relabeled the x-axis of figure 6.1 [the particle model described in abstract] to be the relative frequency of allele A_1 . But that axis differs importantly from the one in figure 6.1 in that it has definite endpoints of 0 and 1. The relative frequency of A_1 cannot go beyond 1 or 0. Furthermore, such boundaries are absorbing boundaries (if we ignore back-mutations), in the sense that, once the population gets to one of those values, it is stuck there. Given a random walk with absorbing boundaries the expectation is that each particle (each population) will eventually move to one of the boundaries. Thus, according to this bit of theory, drift eliminates genetic variation from natural populations. (...) it is not drift *simpliciter* that acts to reduce variation but drift plus the boundaries, which are constraints. (McShea and Brandon 2010, p. 104; see also Brandon and McShea 2012, p. 741)

So, according to these authors, the distribution of probabilities for the movement of a population evolving under drift in the evolutionary space illustrates a *diffusion process*. However once populations reach the limits of the space, they simply accumulate there. This, according to McShea and Brandon, is at least partly due to the fact that the space being considered has some limits that cannot be trespassed.

What matters here is not the adequacy of this response. Rather, it is the fact that, in responding this way, they seem to be conceding that if the long term directional predictions considered before (i.e. the fixation of the most frequent allele) were actually derivable from evolutionary models which incorporate *only* genetic drift, then their thesis of the non-directionality of drift would be mistaken.¹⁵ Their point is, however, that these models are not drift-only models, but that they incorporate other constraints. This way of responding shows that they consider the prediction that an allele whose starting frequency is 0.99 will be fixed in the population (a good/reliable probabilistic directional prediction) as a potential threat to their claims, which also shows that they have this stronger objection in mind.¹⁶

¹⁵ At least in the long-term (i.e. remote generation) sense. It would still be debatable whether in the short-term (i.e. from one generation to the next), a direction could be predicted reliably.

¹⁶ There is a possible alternative reading of McShea and Brandon's response, according to which the absence of absorbing barriers simply points out to the impossibility of making deterministic predictions (i.e. absorbing barriers is what one *would* need in order to make them)—going back, then, to the first thesis. Even if that was really their intention, it does not matter for my purposes. All I contend is that the second thesis is a different possible interpretation of their claims. Furthermore, my argument will concern neither the first nor the second interpretations mentioned so far, but a third one detailed in what follows.

And third, their claim of the non-directionality of drift and its illustration via the particle model could be seen as arguing for an even stronger claim: that drift does not bias populations in any particular direction (remember, this is Sober's original criterion for distinguishing forces), since the probability of the particle ending to the right of where it started (i.e. the allele's increase in frequency) is equal to the probability of it ending to the left (the allele's decrease in frequency). In Table 1, both probabilities equal 0.3125. The image of drift they present seems to indicate that they also have this thesis in mind. They think of a set of populations evolving under drift as a field full of leaves, where the wind blows from variable directions with equal probability. At both ends of the field there are garages with open doors, where the wind can blow leaves in, but not out—in this metaphor, the garages are the absorbing barriers (McShea and Brandon 2010, p. 30). So, if the wind blows from random directions, a particular leaf (population) is equally likely to go in any direction in the next immediate moment, regardless of its initial positioning (provided it is not inside one of the “garages”).

Defenders of the force analogy have not addressed this point. On the contrary, they seem to have conceded to it. For example, Pence claims that “under [genetic drift] (at least on the simplest models [...]) *an allele's frequency is equally likely to increase or decrease at each point in time*” (Pence 2016, p. 11, my emphasis).¹⁷ This idea seems to be even more widespread. Not in the context of this debate, but in a well-known evolution textbook, Futuyma states the following:

Imagine, for example, a population of land snails [...] in which (for the sake of argument) offspring inherit exactly the brown or yellow color of their mothers. Suppose 50 snails of each color inhabit a cow pasture. (The proportion of yellow snails is $p = 0.50$). If 2 yellow and 4 brown snails are stepped on by cows, p will change to 0.511. Since it is unlikely that a snail's color affects the chance of its being squashed by cows, the change might just as well have been the reverse, and indeed, it may well be the reverse in another pasture, or in this pasture in the next generation. *In this random process, the chances of increase or decrease in the proportion of yellow snails are equal in each generation, so the proportion will fluctuate.* (Futuyma 1986, p. 227, emphasis added)

In a recent book, Berwick and Chomsky (2015) equate genetic drift with a “drunkard's walk”, in which a drunk person (a population) takes random steps (changes its frequencies, each step representing a generational transition) in a one-dimensional space (the evolutionary space), at the ends of which there are bars (absorbing barriers). Again, in the middle of the field, the drunk person takes steps in either direction with equal probability and no predetermined direction or bias.¹⁸ Many more examples of this can be provided.

¹⁷ It is interesting to note that the qualification “at least in the simplest models” was added to the final version, but was absent in a previous draft of this essay (see Pence 2012, p. 7). This is not to criticize the author (as his work should be evaluated only by its final version), but I think it illustrates the tendency to think in these terms. Another point worth mentioning is that it is not clear what these “simplest models” are. The simplest model that I can think of (at least the simplest one that is minimally useful or realistic) is the particle model, for which, as I will show, Pence's claim is false.

¹⁸ In a previous writing (Roffe 2016 which curiously utilized the same image) I argued that, to correct that image, one must make the field “tilted” like a gabled roof, so that the drunk person has a greater

What I will argue in the following section is that, even conceding that the second version of the objection—i.e. no reliable directional predictions can be made with drift-only models—succeeds, this third and stronger version doesn't obtain.

The biasing effects of genetic drift

In taking the position described above, defenders of drift's force (and directional) status have made it into a strange factor. Indeed, according to them, drift can bias the direction in which populations go in the long term (in a large number of generations) but not in the short term (for each particular generational transition).¹⁹ In contrast, Newtonian mechanics forces' long term directions follow from their short term ones. That is, usually, if one gets one kind of direction, then one also gets the other, either by adding up or dividing down. How could this be the case?

One possibility would be the following: Consider a game where a fair coin is tossed, heads yields +1 points, tails yields -1 points, and landing sideways yields 0 points. Let us say that the coin is tossed until one reaches either 0 or 100 points (the game ends when one reaches either outcome). If the starting point is at 50 points, then the probability of finishing at 0 or at 100 is the same. But if we start at 80 points, then there is a greater probability of ending up at 100 points than at 0 points, despite the fact that *in every single throw the probability of getting +1 or -1 is the same*. Here the coin would represent the population as a whole—not an individual within the population—and the coin toss a generational transition; obtaining heads represents obtaining a sample in which the allele frequency for an *A* allele is greater than the original population's frequency of it; the reverse occurs with tails. Landing sideways represents the case where the population and sample frequencies coincide (there are no sampling errors). Finally, the number of points represents the position of the population within the evolutionary space. Thus, drift could tend to fix the most frequent alleles simply because one starts out closer to one of the absorbing barriers than the other, so less “steps” are necessary in that direction to become fixed.

What I will argue in this section is that this, and the previous section's images of a population evolving under drift, are all wrong, since the probability that an allele increases in frequency is not the same as the probability that it decreases in frequency, in all cases except when the starting frequencies are 0.5. This simple fact


Footnote 18 continued

chance of walking one way or the other depending on where she is in the space (for the reason, see below).

¹⁹ An example might help clarify what I mean by “short term” and “long term”. Take a haploid population that has a frequency of 0.8 for an allele *A* at generation g_n . Saying that at the limit (after an infinite number of generations have passed) *A* will increase in frequency (it will actually become fixed with probability 0.8) is not the same as saying that *in generation g_{n+1}* (the immediate next generation), *A* will increase in frequency. I label the first a “long term” prediction, since it tells us about what will happen in some (possibly remote) generation; “short term” effects deal with the immediate next generation to the one being considered (for simplicity's sake, I am assuming generations to be discrete and non-overlapping). The distinction is important because many take the first prediction to be derivable from population genetics theory, but not the second one.

Table 2 Probability distribution for a new particle model

s_4	0.0039	0.0469	0.2109	0.4219	0.3164
s_3	0.0156		0.1406	0.4219	0.4219
s_2		0.0625		0.375	0.5625
s_1			0.25		0.75
			1		
0	0.25	0.5	0.75	1	

**Fig. 2** Biological interpretation of Table 2

follows, for example, from the Wright–Fisher model of drift, but it can be illustrated in simpler terms, again, utilizing the particle model.

Let us now consider a population evolving under drift with a starting frequency different than 0.5, say, 0.75 for A and 0.25 for a . Since, as was previously explained (see Sect. 2), the probabilities of the particle going right and left at each step correspond to the probability of sampling an A or an a allele (which, in turn, correspond to their frequencies in the gamete population), the particle model has to be modified to make $P(\text{the particle goes right at the next step}) = 0.75$ and $P(\text{left}) = 0.25$. When this is done, the model looks as in Table 2.

As before, this instance of the particle model should be interpreted biologically in the way indicated by Fig. 2.

Notice that the probability that the frequency of A goes up (or I should say “right”) to 1 is higher than the sum of the probabilities that it goes left (either to 0.5, 0.25 or 0)—the probabilities are 0.3164 and 0.2617 respectively. That is, the claim that *in a drifting population, the probability that an allele frequency increases is equal to the probability that it decreases*²⁰—it is only true for populations whose starting allelic frequencies are 0.5. In contrast, if the starting

²⁰ What is not wrong is the (perhaps confusingly similar) claim that, if there is no selection, mutation, etc., then every *particular* allele has the same probability of being sampled in every generation, but not any allele *type*.

frequency for an allele is above 0.5, the model implies that its frequency will increase in the next generation with a higher probability than it predicts that it will decrease (and vice versa if it is below 0.5).²¹

To avoid confusions as to exactly what I'm claiming here, notice that the *expected value* of the frequency of the *A* allele in the daughter population is the same as its frequency in the parent population (0.75 in the case of Table 2; Fig. 2). I do not claim that the expected outcome is that the frequency increases. Rather, the point is that (when $p > 0.5$) the summed probabilities that the population ends up with a frequency *higher than the expected value* are greater than the summed probabilities that it ends up *lower than the expected value*. To put it in a different way, my discussion concerns the situations where an *unexpected* result occurs. This also does not imply that one can reliably predict that the population will in fact go right, since that only occurs with probability 0.3164 (i.e. it comes out true only slightly more than 3 out of 10 times one makes it). However, even if that is the case, there is a clear sense in which there is a bias operating. In the coin tossing analogy, one gets this by making the coin *biased* towards heads as the number of points surpasses 50 (and vice versa).²²

The idea that drift is generally biased with respect to the frequencies of the allele types is important because, as was said, in Sober's original presentation, the common currency between the evolutionary forces was their "biasing effect" on allele or genotype frequencies. Therefore, this would seem to show that, at least in Sober's original sense, drift could be considered a force.²³ Even McShea and Brandon themselves acknowledge that "it is easy to see that a probabilistic process can make directional predictions (...). A random walk driven by flips of a biased coin is probabilistic and directional. (...) [T]he directionality is possible because of a bias. (McShea and Brandon 2010, p. 14). However, they claim that this is not the case for drift.²⁴ If what I'm saying is correct (and the coin image from above is adequate) then it seems that under their own standards they should consider drift to be a directional factor.

Additional consequences can be drawn from this. For example, it is possible to see that the long term directional effects of drift actually do follow (at least partially) from its short term ones. For a population with a starting frequency of $p = 0.9$, fixing the *A* allele is more likely than fixing the *a* allele, not only because it starts out "closer" to the absorbing barrier to the "right" of the space (i.e. the frequency of 1), but also because it is more likely that in the very next generation the population frequency will increase rather than decrease (it is more likely that the

²¹ Some qualifications must be made here, see the "Mathematical appendix" section.

²² For the way in which the biases must change according to the number of points, see the "Mathematical appendix" section below.

²³ Perhaps there are other reasons for not doing so. I actually find McShea and Brandon's "drift as constitutive" argument more adequate (see Roffé 2016). The point is only that drift shouldn't be discarded as a force on the basis that it can *never* (or almost never) bias populations one way or the other.

²⁴ Or, in the original context of this quote, for the ZFEL: "But the ZFEL arises from an *unbiased* random process, and yet it predicts directional change: increasing diversity and complexity. How is this possible?" (McShea and Brandon 2010, p. 14).

population will go that way). Therefore, for those willing to defend the status of drift as a force, this would imply that it is a less peculiar one than was originally thought.

Lastly, another interesting question concerns the following: I doubt any of the mentioned authors do not know the results presented in this section (none of them would doubt, I imagine, my Table 2). The question is then why these facts haven't been noticed before within the discussion, and why it seems so common to make the kind of mistake mentioned above. I can imagine two reasons. First, because we tend to take as paradigmatic populations whose allelic frequencies are 0.5. McShea and Brandon's presentation of the particle model I think illustrates this pretty well, since it assumes the starting frequencies are precisely those.²⁵ The important point to make here is that what looks as a paradigmatic case is actually the exception, at least regarding the directionality issue, and that one should be careful not to extrapolate from this particular case. Second, because we tend to think in terms of infinite (or very large) populations, for which the sort of bias described above tends to disappear. All of this points to some of the heuristic and educational benefits mentioned above, and which this discussion potentially gives. Not taking those kinds of populations as paradigmatic, at least when discussing drift, has some immediate benefits, especially given the role that paradigmatic cases have in knowledge acquisition (Kuhn 1970) and in the regulation of philosophical discussions.

Some possible objections

This section examines some possible objections to my arguments. I consider two of them. First, I examine more closely the notion of bias, and the ways in which it is legitimate to speak about biases in drift cases. Then, I touch upon the drift as a process versus drift as an outcome discussion and how it affects my position here.

Speaking of biases in drift introduces a conceptual imprecision. It could be objected that (as McShea and Brandon note) a sampling of gametes whose initial frequencies are different than 0.5, and a series of tosses of a biased coin are not conceptually equivalent processes.

From the point of view of probability theory, a sampling process consisting of a series of n tosses of a fair coin assigns the same probability values as a sampling process in which n balls, which are of two types (let's say some marked with an H, and some with a T), and are present in equal quantities, are taken (with replacement) from an urn. The same happens with an unfair coin (for which, for example, the probability of landing heads is 0.75) and a sampling of balls from an urn, where the frequencies of H and T are 0.75 and 0.25, respectively. Nevertheless, in this last case, it is usually said that one of these processes is random while the other one isn't. This is due to the fact that, in probability theory, a random sampling process is defined as one where every point in the sample space (each one of the *particular*

²⁵ Notice also that in Futuyma's quote, the population starts with a frequency of $p = 0.5$, and the problem arises when he extends the result he obtains from the first generation to the second.

possible samples) has the same probability of occurring (see for example Feller 1971, p. 30).

For example, if we take a sample of two individuals from an urn containing four balls,²⁶ $\{H_1, H_2, H_3, T_4\}$, then every *particular* sample will have the same probability of occurring. For instance, the sample $\{H_1, H_2\}$ has the same probability of occurrence as the sample $\{H_1, T_4\}$. The same does not happen with the points of the sample space of the loaded coin experiment. The sample $\{\text{Heads}_1, \text{Heads}_2\}$ has a different probability of occurrence than $\{\text{Heads}_1, \text{Tails}_2\}$. The equivalence between both processes holds only between complex *events*, such as “we obtain an H ball and a T ball” and “we obtain one heads and one tails”.

I concede this point, and I also concede also that—in the technical sense—the sampling processes used to model generational transitions are only non-random when evolutionary factors like natural selection intervene. If, in addition, “biased” is understood as “non-random”, then it is also not adequate to speak of drift as biased.

However, I think it is misleading to confuse the philosophical distinction between directional and non-directional factors with the technical distinction between biased and random processes. To confuse the two would be to lose sight of the point being discussed. Remember that what is being discussed is if, for a population evolving under drift, an increase (or a decrease) in the frequency of an allele can be predicted a priori. The fact that a population evolving under drift alone tends to evolve more towards the increase than towards the decrease of such frequency is, at any rate, a directional effect. I believe that what Sober was calling “biasing effects” is something closer to this than to the technical notion of bias. A random process (in the technical sense) can have directional effects even *at the same level*.²⁷

A second objection that could be leveled against my argument is that I only show that the *outcomes* of drift are biased, not that it is biased as a *process*. Since I’m speaking about drift only as it is modeled on theories like population genetics (because the relevant issue at hand is the structural similarity between two *theories*), it is not correct then to speak of drift as a process. The only process involved is the sampling of parents or gametes, which can be affected either by selection, mutation, etc. and in which sampling errors can occur. Strictly speaking, it is this process (the sampling of gametes) which I argue is biased (in the non-technical sense), when populations are finite, selection, mutation, etc., are absent and initial frequencies are different than 0.5. Perhaps there is an interesting sense in which one can speak of drift as a process, but I don’t believe it to be relevant to the discussion at hand. The present discussion is orthogonal to the “drift as a process versus drift as an outcome” one.

²⁶ Alternatively, this can be thought of as taking a sample of two alleles from a gamete population of 4 alleles.

²⁷ Unlike McShea and Brandon’s ZFEL (see their 2010), which predicts these effects only at the next level (the level of ensembles of populations).

Conclusion

In this article I have examined the issue of whether drift can be said to be a directional evolutionary factor. I have argued that this question can be analyzed in (at least) two different ways: whether one can reliably predict, in an a priori manner, the direction in which a population evolving under drift will go towards; and whether drift can bias populations towards one of those directions. The preceding discussion shows that, even if one concedes that the answer to the first question is negative, the answer to the second is positive. In other words, when an allele frequency is greater than 0.5, it is more likely that this frequency will go up rather than the reverse, even in the very next generation (and vice versa when this frequency is lower than 0.5—for more on this see the “[Mathematical appendix](#)” section).

The remaining question is what this discussion implies for the force analogy, and for the status of drift as a force within it. As was said, the question is not whether drift *really is* a Newtonian force. Taken literally, that claim is obviously false. Taken metaphorically, it cannot be said to be true or false in any absolute sense. If newtonian mechanics and evolutionary theory (and their respective “forces”) are analogous in any interesting sense, then they are not identical (an analogy between two identical things would be trivial and valueless), and thus we should expect there to be some differences between them. That is, we should only expect them to be alike in some respects, and not in others. If drift is taken to be a force, and the deterministic or probabilistic character of the theories’ respective “forces” is considered, then they are not alike.²⁸ If the possibility of predicting a direction in an a priori manner is considered, then it is doubtful whether they are analogous or not (they seem not to be in the short term, but they may be in the long term). My point is only that if their forces’ “biasing effect” on the “spatial location” of the entities to which they apply is considered, then they are alike, in a relevant and previously unnoticed sense.

What does this imply about the status of drift as a force? Which of these criteria should be privileged in laying out the force analogy? Metatheoretically, it is important to point out every one of these things, to recognize where the analogy can be successful as well as what its limitations are, in order to have an accurate picture of the extent to which newtonian mechanics and population genetics employ similar explanatory strategies. So, within that discussion, my position would be that drift is like a Newtonian force in some respects, and not in others. Within the pragmatic and educational context (should we remove the analogy from textbooks?), whether the disanalogies are enough to make force language heuristically inadequate must be evaluated on other (i.e. pragmatic or even empirical) grounds. The project of identifying what circumstances make a force explanation heuristically useful in general (and in this case, in particular) cries out to be taken, but is not the object of this paper. However, I do believe that if my points are taken into consideration, the discussion moving forward will be a more fruitful and adequate one. The points

²⁸ Or, again, as Pence (2016) argues with brownian motion as an example, perhaps they are. But this is besides the point here, I am conceding this for the sake of argument.

addressed can also have consequences independently of the force analogy issue, since—as was shown—treating drift as a directionless (non-biased) factor is a more widespread phenomenon. The present discussion shows we should be more careful to avoid the kind of images and extrapolations from particular situations when speaking and thinking about genetic drift.

Acknowledgements This work has been funded by the research projects PICT-2014 No. 1741 (ANPCyT, Argentina), PUNQ 1401/15 (National University of Quilmes, Argentina) and UNTREF 32/15 255 (Universidad Tres de Febrero, Argentina). I wish to thank the members of the philosophy of biology group directed by Santiago Ginnobili; I also thank Charles Pence, and an anonymous referee for helpful comments and criticisms on previous versions.

Appendix

Mathematical appendix

In this appendix, a number of precisions are introduced regarding the directional effects mentioned above. I first distinguish between general and particular effects (predictable in every case, or only in a limited subset). The second subsection analyzes if the biases mentioned can be said to have any general direction. And lastly, in the third subsection, McShea and Brandon's idea that drift's magnitude can be predicted while its direction cannot is revised.

General and particular directional effects

It is clear that, in each particular case, given enough information, we can predict the probability that a population goes “right” or “left” (i.e. an allele increases or decreases in frequency) in the evolutionary space. But, the discussion above suggests much more: that we can predict that, *in general*, drift will bias populations to move towards an increase of the most frequent allele, in every generation. This also seems like an intuitive result: If a population with two kinds of things has more of one, then we should expect deviations of the expected value to occur in that one's direction. The next subsection examines in greater detail whether drift can be said to have this general directional effects.

General directional effects

This subsection examines the generality of the claim that, in every generation, drift will tend to move populations towards an increase of the most frequent allele more than the reverse. Unfortunately, for various reasons, this is not true in general. First, while we can treat gamete sampling as being *with replacement*, parental sampling is a process *without replacement*. Now, for this second type of sampling, there are counterexamples to the claim made above. For instance, consider a haploid population with two alleles (A and a) in one *locus*. If population size at birth is $N = 9$, the frequency of allele A equals $2/3$ (there are 6 A individuals, and so A is the most frequent allele), and the sample size (adult population size) is $n = 6$, then

the expected number of A individuals in the adult population is 4. However, the probability that allele frequency increases (that is, that the sample contains either 5 or 6 A individuals) is smaller than the probability that it decreases (that it contains all 3 a individuals in it)—the probabilities are 0.226 and 0.238 respectively.

The question that arises is, then, if the result holds for sampling processes *with* replacement (like gamete sampling, or parental sampling when populations are very large). The answer is that this also isn't the case. Take $p = 0.51$ and $n = 2$, for example. What is happening here is that the expected value is not a possible sample value, so a frequency of 0.5 in the sample (i.e. when the sample contains one individual of each kind) will count as a case where the frequency of A decreased. This problem arises because the mean is close to, but not exactly, 2. It might be argued that this is merely an effect of the "incommensurability" between the population and sample sizes, an artifact of the choice of sample size, and that it doesn't reflect drift's "intrinsic" direction.

These kinds of cases can be eliminated by putting some mathematical restrictions in place, for example, by establishing that $np \in \mathbb{Z}$. With this restriction, the problem of the general direction can now be analyzed as follows. Call \hat{p} the sample frequency of allele A . What we would need to establish mathematically to prove our general claim is that "If $p > 0.5$, then $P(\hat{p} > p) > P(\hat{p} < p)$ ". This is, of course, equivalent to claiming that "If $p > 0.5$, then $P(\hat{p} \geq p) > P(\hat{p} \leq p)$ ". In a sampling process without replacement, this translates to:

$$\sum_{k=np}^n \binom{n}{k} p^k (1-p)^{n-k} > \sum_{k=0}^{np} \binom{n}{k} p^k (1-p)^{n-k}$$

Even though this would be an interesting result, even from a purely mathematical point of view, I haven't found any mention of it in the mathematical (or biological) literature. And, probably due to my limited abilities in mathematics, I haven't been able to prove it (nor refute it) myself, at least not generally. However, what I have been able to do is to get certain partial results that I think may help illuminate how drift is working in these kind of models.

For example, it is possible to show²⁹ that the result holds for the case where $p = 1 - (1/n)$. This case is interesting because it is one of the most counter-intuitive ones. That is, if in a population of 1000 alleles (gametes) there are 990 A ones, and a sample with replacement is taken of 100 individuals, then what the result claims is that the probability of obtaining 100 A individuals is higher than *the sum* of the probabilities that 1, 2, 3, ..., 98 A alleles are obtained.

Theorem 1 *If $p = 1 - \frac{1}{n}$, then $P(\hat{p} > p) > P(\hat{p} < p)$*

Formally, what needs to be established is that:

$$\binom{n}{n} p^n > \sum_{k=0}^{n-2} \binom{n}{k} p^k (1-p)^{n-k},$$

²⁹ I thank Bjørn Kjos-Hanssen for his help with this part of the proof.

Since $\binom{n}{n} = 1$, and replacing p for the assumed value, this is equivalent to:

$$\left(1 - \frac{1}{n}\right)^n > \sum_{k=0}^{n-2} \binom{n}{k} \left(1 - \frac{1}{n}\right)^k \left(\frac{1}{n}\right)^{n-k},$$

Multiplying both sides by n^n , we get

$$(n - 1)^n > \sum_{k=0}^{n-2} \binom{n}{k} (n - 1)^k.$$

Next, the following lemma can be established:

Lemma 1

$$\sum_{k=0}^{n-2} (n - 1)^k \binom{n}{k} = n^n - (n - 1)^n - n(n - 1)^{n-1}$$

Proof of Lemma 1

$$\sum_{k=0}^{n-2} (n - 1)^k \binom{n}{k} = \left(\sum_{k=0}^n (n - 1)^k \binom{n}{k}\right) - (n - 1)^n \binom{n}{n} - (n - 1)^{n-1} \binom{n}{n - 1}$$

The first term in the subtraction can be operated with as follows. Given the binomial theorem, which claims that:

$$(1 + t)^a = \binom{a}{0}t^0 + \binom{a}{1}t^1 + \binom{a}{2}t^2 + \binom{a}{3}t^3$$

The first term reduces to:

$$\begin{aligned} \sum_{k=0}^n (n - 1)^k \binom{n}{k} &= \binom{n}{0}(n - 1)^0 + \binom{n}{1}(n - 1)^1 + \binom{n}{2}(n - 1)^2 + \dots + \binom{n}{n}(n - 1)^n \\ &= (1 + (n - 1))^n = n^n \end{aligned}$$

The third term can also be simplified as follows, which directly gives us the desired result:

$$\binom{n}{n - 1} = \frac{n!}{(n - 1)!(n - (n - 1))!} = \frac{n!}{(n - 1)!1!} = n$$

□

Given this lemma, all that remains to be proved is that $(n - 1)^n > n^n - (n - 1)^n - n(n - 1)^{n-1}$, that is, that $n^n < 2(n - 1)^n + n(n - 1)^{n-1}$.

Multiplying both sides by $\frac{1}{(n-1)^n}$, we get

$$\left(1 + \frac{1}{n-1}\right)^n < 2 + \frac{n}{n-1} = 2 + \frac{(n-1)+1}{n-1} = 3 + \frac{1}{n-1}$$

or with $m = n - 1$,

$$\left(1 + \frac{1}{m}\right)^{m+1} < 3 + \frac{1}{m},$$

$$\left(1 + \frac{1}{m}\right)^m < \frac{3 + \frac{1}{m}}{1 + \frac{1}{m}}$$

Multiplying the term on the right by $\frac{m}{m}$

$$= \frac{3m + 1}{m + 1} = \frac{3m + 3 - 2}{m + 1} = \frac{3(m + 1) - 2}{m + 1} = \frac{3(m + 1)}{m + 1} - \frac{2}{m + 1} = 3 - \frac{2}{m + 1}$$

What is left to prove is that:

$$\left(1 + \frac{1}{m}\right)^m < 3 - \frac{2}{m + 1}$$

And this is true, since it can be checked directly for $m \in \{2, \dots, 6\}$, and the left side is bounded by e , while the right side (which represents an increasing function) surpasses e by $m = 7$. □

A possible demonstration strategy of the general result would be the following. We first show that, for a starting frequency of $p = 0.5$, $P(\hat{p} > p) = P(\hat{p} < p)$. This is fairly easy to establish. We need to show that:

Theorem 2

$$\sum_{k=0}^{n/2} \binom{n}{k} \left(\frac{1}{2}\right)^k \left(\frac{1}{2}\right)^{n-k} = \sum_{k=n/2}^n \binom{n}{k} \left(\frac{1}{2}\right)^k \left(\frac{1}{2}\right)^{n-k}$$

Proof of Theorem 2 The strategy here is to “pair up” the factors of the sums, such that the probability that $k = 0$ on the left “pairs” with $k = n$ on the right; $k = 1$ on the left with $k = n - 1$ on the right... and, in general, $k = r$ in the left with $k = n - r$ on the right (Fig. 3).

Since $\left(\frac{1}{2}\right)^{n-k} \left(\frac{1}{2}\right)^{n-(n-k)} = \left(\frac{1}{2}\right)^k \left(\frac{1}{2}\right)^{n-k}$, and since




Fig. 3 Pairing of the factors of the sums

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} = \frac{n!}{(n-(n-k))!(n-k)!} = \binom{n}{n-k}$$

then each of the paired factors will be equal to the other. □

Next, I establish the following conjecture: as p increases by “steps” of $1/n$, $P(\hat{p} \geq p)$ increases as well. That would give us a sort of induction over p (theorem 2 being the base step, and this conjecture the inductive step). Formally, we need that:

$$\binom{n}{n} \left(1 - \frac{1}{n}\right)^n \left(\frac{1}{n}\right)^0 > \binom{n}{n-1} \left(1 - \frac{2}{n}\right)^n \left(\frac{2}{n}\right)^0 + \binom{n}{n-1} \left(1 - \frac{2}{n}\right)^{n-1} \left(\frac{1}{n}\right)^1 > \dots$$

That is, that

Conjecture 1 $\sum_{k=0}^{r+1} \binom{n}{n-k} \left(1 - \frac{r+1}{n}\right)^{n-k} \left(\frac{r+1}{n}\right)^k$ is a decreasing function for $r = 0, 1, 2, \dots, n - 1$.

If this could be established it would seem to show something else. Put in the terminology of the coin game image, it would seem to show that the bias of the coin gets bigger and bigger the more one gets closer to one of the absorbing states. However, that would be partially misleading, since proving that $P(\hat{p} \geq p)$ increases when p increases by steps of $1/n$ is not equivalent to proving that $P(\hat{p} > p)$ increases in the same way. This last statement is, in fact, false. Consider the case where $n = 10$; for $p = 0.9$, $P(\hat{p} > p) = 0.348$, while for $p = 0.8$, $P(\hat{p} > p) = 0.375$. It does hold however, in this case, for $P(\hat{p} \geq p)$. This is due to the fact that the increase in the probability of the expected value ($\hat{p} = p$) is greater than the decrease of the probability of going right.

If the conjecture holds, then this would paint an interesting and complicated picture for the changes in the coins’ bias. That is, the bias would still exist (and the main result needed to claim that drift has a general direction would be established), but the changes in the bias would be somewhat strange: as one moves closer to 100 points, the probability of landing heads or sideways increases. However, at the same time, it could happen that the coin gets closer to 100 points and the bias gets smaller, while this decrease would be compensated by the probability of it landing sideways getting bigger.

In any case, two conclusions from the preceding discussion can be drawn. First, that whether or not drift has a general direction, particular directions can be predicted a priori in particular cases, and this is enough to make the point of the article. And second, that if the conjecture holds, then the general claim can be made, and it is only for the changes in the bias that one gets more complicated results.

The magnitude of drift

Lastly, another interesting point concerns notion of the magnitude of drift, and its differences with the direction. It is generally assumed—e.g. it was by Brandon

2006, p. 324—that the magnitude of drift can be predicted a priori, in a sense that its direction cannot, by considering population sizes.

What the previous considerations show is that the expected *magnitude* of drift depends not only on the population size, but also on the starting frequencies of the alleles. The last example shows a case where increasing the frequency of an *A* allele resulted in an increase of $P(\hat{p} = p)$ and in a decrease of both $P(\hat{p} < p)$ and $P(\hat{p} > p)$. That is, a mere increase in the frequency of an allele, without modifying population sizes, made drift “weaker”—it diminished the probability that the sample values deviate from the expected ones. In fact, both magnitude and direction depend upon both factors, population size and initial frequencies.

Additionally, as it happens with the case of the direction, it is not possible to obtain a deterministic prediction about the magnitude of drift simply by specifying the population sizes and the initial frequencies (e.g. a prediction that tells us that *p* will either increase or decrease by 0.1). At most, specifying that would allow us to say, with a given degree of confidence, that the sample values will occur within a certain interval (the smaller the interval, the less confidence we have). But this prediction is also different than the one given by selection only models—it doesn’t state a particular magnitude, but a range of magnitudes. In sum, it seems that whatever is said about drift’s direction, must also be said about its magnitude.

References

- Barrett M, Clatterbuck H, Goldsby M et al (2012) Puzzles for ZFEL, McShea and Brandon’s zero force evolutionary law. *Biol Philos* 27:723–735
- Beatty J (1984) Chance and natural selection. *Philos Sci* 51:183–211
- Berwick RC, Chomsky N (2015) Why only us: language and evolution. MIT Press, Cambridge
- Brandon RN (2006) The principle of drift: biology’s first law. *J Philos* 103:319–335
- Brandon RN, McShea DW (2012) Four solutions for four puzzles. *Biol Philos* 27:737–744
- Crow JF, Kimura M (1970) An introduction to population genetics theory. The Blackburn Press, New Jersey
- Darwin C (1872) On the origin of species by means of natural selection, 6th edn. John Murray, London
- Díez J, Lorenzano P (2013) Who got what wrong? Fodor and Piattelli on Darwin: guiding principles and explanatory models in natural selection. *Erkenntnis* 78:1143–1175
- Earnshaw E (2015) Evolutionary forces and the Hardy–Weinberg equilibrium. *Biol Philos* 30:423–437
- Feller W (1971) An introduction to probability theory and its applications. Wiley, New York
- Futuyma DJ (1986) Evolutionary biology, 2nd edn. Sinauer Associates, Sunderland
- Gillespie JH (1998) Population genetics: a concise guide. Johns Hopkins University Press, Baltimore
- Hartl DL, Clark AG (2007) Principles of population genetics, 4th edn. Sinauer Associates, Sunderland
- Hitchcock C, Velasco JD (2014) Evolutionary and Newtonian forces. *Open Access J Philos* 1:39–77
- Kuhn TS (1970) The structure of scientific revolutions, 2nd edn. The University of Chicago Press, Chicago
- Matthen M, Ariew A (2002) Two ways of thinking about fitness and natural selection. *J Philos* 99:55–83
- McShea D, Brandon R (2010) Biology’s first law. The University of Chicago Press, Chicago
- Pence CH (2012) It’s okay to call genetic drift a “Force”. Available at: <http://philsci-archiv.pitt.edu/9256/>
- Pence CH (2016) Is genetic drift a force? *Synthese*. doi:10.1007/s11229-016-1031-2
- Roffe AJ (2016) La genética de poblaciones y las fuerzas evolutivas. Undergraduate dissertation, Universidad de Buenos Aires

- Roffe AJ, Ginnobili S (in press) ¿Son los genetistas de poblaciones inductivistas estrechos? *Scientiae Studia* 15:2
- Rudner R (1953) Value judgments in the acceptance of theories. In: Frank P (ed) *The validation of scientific theories*. The Beacon Press, Boston, pp 24–28
- Skipper RA (2006) Stochastic evolutionary dynamics: drift versus draft. *Philos Sci* 73:655–665
- Sober E (1984) *The nature of selection: evolutionary theory in philosophical focus*. University of Chicago Press, Chicago
- Sober E (2011) A priori causal models of natural selection. *Australas J Philos* 89:571–589
- Stephens C (2004) Selection, drift, and the “forces” of evolution. *Philos Sci* 71:550–570
- Williams MB (1980) Similarities and differences between evolutionary theory and the theories of physics. *PSA Proc Bienn Meet Philos Sci Assoc* 2:385–396