

Four Pillars of Statisticalism

Denis M. Walsh,* André Ariew,[†] Mohan Matthen[‡]

Over the past fifteen years there has been a considerable amount of debate concerning what theoretical population dynamic models tell us about the nature of natural selection and drift. On the causal interpretation, these models describe the causes of population change. On the statistical interpretation the models of population dynamics models specify statistical parameters that explain, predict, and quantify changes in population structure, *without* identifying the causes of those changes. Selection and drift are part of a statistical description of population change; they are not discrete, apportionable causes. Our objective here is to provide a definitive statement of the statistical position, so as to allay some confusions in the current literature. We outline four commitments that are central to statisticalism; They are: 1. Natural Selection is a higher order effect; 2. Trait fitness is primitive; 3. Modern Synthesis (MS)-models are substrate neutral; 4. MS-selection and drift are model-relative.

Keywords

keyword1 • keyword2

0 Introduction

Models of evolutionary population dynamics represent changes in the structure of evolving populations in terms of selection, drift, and fitness.¹ There is a range of such models, from the Wright-Fisher equations of population genetics, to those of quantitative genetics, to the Price and breeders' equations. They vary in their structure and the concepts they deploy, but the common denominator is that in all of them fitness explains and predicts change in population structure. A population that varies in fitness is said to be undergoing selection. When the actual amount of change in such a population differs from the expected change, the population is said to manifest drift. Over the past fifteen years there has been a considerable amount of debate about what theoretical population dynamics models tell us about the biological reality of selection and drift. Two major positions have emerged, the causal and the statistical.

The orthodox, causal approach recommends that these models should be interpreted as descriptions of the causes of population change. On this view, natural selection and drift are forces, or causal processes,

¹There are other explanatory factors, including mutation and migration. For convenience, we restrict our attention to selection and drift here, in order to concentrate on the issues raised in the debate so far, but we recognize that there are important things to be said about migration and mutation. We thank an anonymous referee for raising this point.

*Department of Philosophy, Institute for the History and Philosophy of Science and Technology, University of Toronto, 170 St. George Street, Toronto, ON, Canada M5R 2M8, denis.walsh@utoronto.ca

[†]Department of Philosophy, University of Missouri, 438 Strickland Hall, Columbia, MO 65211-4160 USA, ariewa@missouri.edu

[‡]Department of Philosophy, University of Toronto, 170 St. George Street, Toronto, ON, Canada M5R 2M8, mohan.matthen@utoronto.ca

Received 9 September 2016; Accepted 9 November 2016

DOI: TBD



that propel a population through changes in trait frequencies (Sober 1993). Variation in trait fitness is a measure of the strength of selection. The degree of divergence from the expected change in trait frequencies measures the efficacy of drift.

Selection and drift, on this view are independent and proprietary population-level causes of population change (Sober 1993, Stephens 2004). By ‘independent’ we mean that the minimal conditions sufficient for selection to occur can obtain without those that are sufficient for drift (and *vice versa*). They are proprietary in the sense that, each process—selection and drift—has its own distinct kind of effect, each called after its supposed cause. There is thus a distinction to be made between selection the process (which is the cause) and selection the product (which is the effect), and similarly between drift the process and drift the product (Millstein 2002, Stephens 2004). The important point is that on this orthodox causal approach, the population models that quantify the degree of selection and drift in a population do so by articulating and differentiating the population-level causes of evolutionary change.²

The statistical interpretation cautions against reading population models in this way. According to statisticalists, the models of population dynamics provide us with a set of statistical parameters that explain, predict, and quantify changes in population structure, but they do not do so by representing selection and drift as discrete, separable causal processes. Selection is a measure of how much change is to be expected in the frequency of a gene or a trait; drift is a measure of the divergence of the actual change from the expectation.

While the debate between the orthodox and statistical factions has been vigorous, it has not always been particularly productive or germane. This is due, in some measure, to a widespread misapprehensions of the statisticalist position. This, in turn, may arise from the fact that the core commitments of the statisticalist position have never been fully assembled and articulated. Rather, the precepts of the statisticalist view are spread throughout an array of papers by various (combinations of) authors. And, as with its causalist counterpart, the various presentations of statisticalism may encompass a range of minor variations. Our objective here is to state as clearly and simply as possible what we take to be the shared core commitments of the statistical interpretation. First, some context is needed. We offer some background to statisticalism in sections 1 and 2, before proceeding to an articulation of its central tenets in Section 3.

1 Darwinian and Modern Synthesis Selection

Organisms live, die, reproduce, and pass on their heritable traits to their offspring (sometimes in mutated form); they immigrate and emigrate. Darwin’s great insight was that in the aggregate these events are sufficient to account for the fit and diversity of organic form. Darwin reasoned that as a consequence of all this living, reproducing, and (imperfect) inheriting of traits, populations change in their structure. In fact, populations undergo a number of different kinds of changes. They change in their lineage structure; as some organisms produce more offspring than others, their lineages increase their representation in the population. Furthermore, populations undergo changes in their trait distribution; some trait types increase in their relative frequency with respect to others. Additionally, populations change in their adaptedness; they come to comprise individuals in general better suited to surviving and reproducing in their conditions of existence.³ Quite frequently, these changes march in step, which has led to the

²There are variants on this. In one version of the orthodox approach natural selection and drift are said to be two vector forces of evolutionary change (Sober 1993, Millstein 2006, Hitchcock and Velasco 2014). A further variant contends that natural selection is a force but drift is not (Brandon 2006).

³In what follows we will concentrate on the relation between change in trait distribution and change in lineage structure. Walsh, Lewens and Ariew (2002) and Walsh (2003) demonstrate that these are not the same process. There can be change in the lineage structure of a population that is explained and predicted by variation in individual fitness, even when there is no change in trait distribution (and no variation in trait fitness). Earnshaw-Whyte (2012) demonstrates that predictable change in trait distribution can occur without variation in individual survival and reproduction. Kaufmann (2003) demonstrates that change in either lineage or trait distribution under natural selection does not necessarily lead to change in adaptedness.

tacit supposition that they are the same process and are explained in the same way. But this is a mistake. Changes in lineage structure and in trait distribution require (or use) different kinds of explanatory models.

In order to explain or predict change in *lineage* structure from one time to the next, we need to know the propensity of *individuals* to survive and reproduce.⁴ In order to recognize, quantify and predict a population's change in *trait* distribution, we need to know something else. We need to be able to divide the population into abstract trait classes (or 'trait types'). We do this, for each heritable trait, by collecting together all those individuals that share that trait. The population of trait classes cross-classifies the population of individuals—each trait class has many individuals as members, and each individual is a member of many trait classes. Having constructed these abstract classes—trait types—we need to assign to them a parameter that will enable us to account for the population's change in trait distribution. The parameter represents the predicted growth rate of the trait type in the ensemble (Matthen and Ariew 2002). Only with knowledge of these growth rates can we predict and explain the change in a population's trait distribution.

The information required to explain change in trait distribution is more or less irrelevant to the understanding of change in lineage structure and *vice versa*. Given information about the propensities of individuals to survive and reproduce, we do not need to know anything about which trait classes they fall into in order to predict and explain the increase in prevalence of one lineage over another. Conversely, if we know the relative growth rates of the abstract trait types, we need no further information about the propensities of individual organisms in order to account for the change in trait distribution (Walsh, Lewens and Ariew 2002; Walsh 2003). We don't have to know which individuals belong to which trait classes, the ecological conditions, which particular individuals survive and reproduce, or who begets whom. Nor is the minimum information sufficient to explain change in trait distribution sufficient to explain change in lineage structure. Changes in lineage structure and changes in trait distribution are thus different kinds of changes that result from the aggregated activities of organisms in a population. As such, they require different models, with different theoretical parameters.

Unfortunately, the corresponding theoretical parameters of these different models share a name: 'fitness.' 'Fitness' is a crucial theoretical term in evolutionary theory, but an ambiguous one. It can stand variously for a causal property (a propensity) of a concrete entity, an organism. Or it can stand for the growth rate of an abstract trait type. In order to disambiguate the term, we'll call them 'vernacular fitness' and 'trait fitness' respectively.⁵

The disambiguation is important for an obvious reason. It has become a commonplace to characterize natural selection as that process that occurs in a population when there is variation in fitness (Lewontin 1974; Godfrey-Smith 2007). If 'fitness' is polysemous, then 'variation in fitness' also has no univocal meaning. If, in turn, 'natural selection' is defined in terms of variation in fitness, then there should be a corresponding ambiguity in the term 'natural selection'; and there is. The process of selection that Darwin postulated is in essence the change in *lineage structure* that occurs when there is variation in *vernacular fitness*.⁶ Accordingly, we shall call this phenomenon 'Darwinian selection' (or 'D-selection'). The other process, in which populations change in their *trait distribution* as a function of variation in their *trait fitnesses*, we'll call 'Modern Synthesis selection' (or 'MS-selection').

An evolutionary model is a description that represents the dynamics of a population by citing variations in fitness. So, if there are two kinds of fitness, there ought to be two kinds of evolutionary *models*; and there are. One kind represents changes in the lineage structure of a population as a function of variation in vernacular fitness. These models articulate the individual-level, 'ecological' causes of population change. The other kind represents changes in trait distribution as a function of variations in trait fitness.

⁴This is what Matthen and Ariew (2002) call 'vernacular' fitness.

⁵Matthen and Ariew (2002) use the terms 'vernacular' fitness and 'predictive' fitness to mark the same distinction.

⁶He further stipulated, in effect, that when lineages vary with respect to the trait classes they are members of, the population will change in its trait distribution too: "... any variation, ... if it be in any degree profitable to an individual of any species, ... will tend to the preservation of that individual, and will generally be inherited by its offspring" (1872: 61).

In keeping with our conventions, we'll call the former 'D-models' and the latter 'MS-models', to mark the fact that they employ the statistical concepts pioneered by Fisher and others at the inception of the Modern Synthesis.

MS-selection models arose out of the realization that D-selection models of evolution are incomplete. By the early 20th century, evolutionary biologists understood that one cannot explain and predict the magnitude and direction of evolutionary change in trait distribution solely from the survival and reproduction of individuals. We must know the rates of change in relative frequency of the trait types. For that, extra information is typically required, including the effects of inheritance, details of reproductive schedules, mutation and migration rates, demographic factors such as the size and growth rate of the population as a whole, and the composition of variation. MS-selection models were constructed to meet these technical challenges. The objective of these MS-models is to analyze and predict population-level trends in trait frequencies (Nowak 2006, Ariew and Lewontin 2004). Through the work of seminal investigators including RA Fisher, Sewell Wright, and JBS Haldane, key concepts of evolution, including selection, mutation, variation, were embedded in the mathematical framework of MS-models.

A number of authors evince an understanding of the difference between explaining changes in trait distribution by employing MS-models, and citing the individual-level causes of those changes. Wade and Kalizs (1990) distinguish between 'selection models'—those that posit an association between change in mean character value and selection gradient—on one hand, and causal studies of selection on the other. The distinction is also implicit in Sober's (1993) influential discussion of evolutionary 'source laws' and 'consequence laws.'

Whereas it is mainly ecology that tries to provide source laws for natural selection, the consequence laws concerning natural selection are preeminently part of the province of population genetics. It doesn't matter to the equations in population genetics *why* a given population is characterized by a set of selection coefficients ... These values may just as well have dropped out of the sky. (Sober 1993: 59)

Similarly, Okasha gestures toward the difference between the models of population dynamics and the explanation of the ecological causes of evolutionary change.

Population genetics models are (deliberately) silent about the *causes* of the fitness differences between genotypes whose consequences they model. [...] To fully understand evolution, the ecological factors that lead to these fitness differences must also be understood. (Okasha 2015; emphasis in original) ⁷

The important point is that explaining changes in trait distribution and explaining the ecological causes of population change are different projects. They call for very different kinds of models: MS-models and D-models respectively.

Much of the confusion in the debate between the causal and statistical interpretations stems from the failure to appreciate the distinction between D-models and MS-models. Statisticalism is a thesis about the interpretation of MS-models. It has nothing to say about D-models. With the distinction between D-models and MS-models to hand, we can now make a first—crucial—clarification: *The debate between causalists and statisticalists, properly construed, is exclusively about the interpretation of MS-models.*

2 Statisticalism

Statisticalism encompasses two positive claims and a negative claim.

Positive Claim #1: MS-models cite statistical properties of trait types

Positive Claim #2: MS-models explain changes in trait distribution

⁷We thank Fermín Fulda for drawing these passages to our attention. See Fulda ([forthcoming](#)). We also thank an anonymous referee for pointing to the commonalities.

Negative Claim: MS-models do not cite the causes of population change⁸

A strong strain of statisticalism can be found in the writings of Levins and Lewontin (1985) and Marjorie Grene (1961). Grene is explicit: she insists that “... we must ...distinguish between ‘genetical selection,’ which is purely statistical, and Darwinian selection which is environment-based and causal. They remain two distinct concepts with a common name.” (Grene 1961: 31).

The distinction that Grene highlights is as follows. In developing his theory of natural selection, Darwin appeals to certain causal intuitions about the consequences for a population of individuals succeeding differently in the struggle for existence. In the first four chapters of *The Origin*, Darwin presents a series of thought experiments designed to convince us that variation in organisms’ ability to succeed in the struggle for existence can lead to population change (Lennox 1991). In developing the ‘Genetical Theory of Natural Selection’, Fisher (1930) doesn’t appeal to causal intuitions at all. Instead he draws upon statistical intuitions (which are, admittedly, harder to come by) concerning what should happen to a population in which there is additive variance in the growth rates of its abstract types. Margaret Morrison nicely captures the core of Fisher’s statistical method:

Essentially, he treated large numbers of genes in a similar way to the treatment of large numbers of molecules and atoms in statistical mechanics. By making these simplifying and idealizing assumptions, Fisher was able to calculate statistical averages that applied to populations of genes in a way *analogous* to calculating the behavior of molecules that constitute a gas. (Morrison 2002: 58-59 emphasis in original)

Where Darwin characterizes evolution as an aggregate causal consequence of the actions of individual organisms (“all these things follow inevitably from the struggle for existence”), Fisher describes evolutionary population change as an analytic consequence of the statistical structure of a population of abstract entities, trait types (or, as he called them, ‘gene ratios’). For instance, his fundamental theorem says: “The rate of increase of fitness of any [population] at any time is equal to the additive genetic variance at that time” (Fisher 1930, 36). Again, Grene is keenly aware of the implications: “In terms of the ‘strictly defined’ concepts it uses, the fundamental theorem is not a theorem of natural selection, but a statistical device for recording and predicting population changes” (1961: 30).

Grene’s distinction between ‘genetical selection’ and ‘Darwinian selection’ provides an historical precedent for statisticalism. She advocates Positive Claim #1 and the Negative Claim (above), but pre-scinds from Positive Claim #2. Unlike Grene, contemporary statisticalists also argue for Positive Claim #2—while MS-models do not identify the causes of population change, they are nevertheless genuinely explanatory (Matthen and Ariew 2002, Walsh Lewens and Ariew 2002). Statisticalists are thus committed to a class of statistical, non-causal explanations (Ariew et al 2014, Walsh 2013, 2015, Matthen and Ariew 2009).⁹ The plausibility of statisticalism rests on its ability to demonstrate that while change in trait distribution is caused by the capacities of individual organisms to survive and reproduce, it is nevertheless explained (and *not* caused) by the statistical properties (fitnesses) of abstract trait types. The task of promoting statisticalism is the task of making Positive Claim #1, Positive Claim #2, and the Negative Claim more palatable. Statisticalists have offered a variety of different arguments each of which is aimed at supporting some subset of the three claims.¹⁰ More importantly, behind each of the arguments there are four core commitments, which as yet have gone unarticulated.

3 Four Core Commitments

The four core commitments of statisticalism are:

⁸To fill this out a little more: MS-models do not cite the causes of population change, which are all events involving individual organisms. These models aggregate the individual level causes of population change, and appeal to their statistical properties.

⁹Note statisticalists do not deny that any instance of evolutionary change involves causal events. They do deny that MS-selection explanations appeal to them.

¹⁰It is not our objective to rehearse them all here.

- i. Natural selection is a higher order effect.
- ii. Trait fitness is a primitive concept.
- iii. MS-models are substrate neutral
- iv. MS-selection and drift are model-relative, that is to say a population can be said to be undergoing MS-selection and/or drift only relative to an MS-model.

These are the four pillars on which the statistical interpretation rests.¹¹

Natural selection is a higher order effect.

There are basically just two kinds of change to individuals that can bring about evolutionary changes in the trait distribution of a population (Earnshawe-Whyte 2012). Organisms may leave a population, or they may enter. (We leave aside for now the changes in trait distribution that may be initiated by individuals changing their traits, through mutation, non-Mendelian segregation, or lateral gene transfer.) Organisms leave a population by emigrating or by dying. Organisms enter a population by immigrating or by being born (hatched, germinated, budded—whatever). Systematic changes in the trait distribution of a population occur when the entry and departure of individuals is biased. Roughly speaking, when disproportionately more individuals possessing traits of one type leave or enter than individuals possessing traits of another type, a change in the population's trait distribution ensues.

To amplify: evolutionary changes in population trait distribution are *analytic* consequences of the activities of organisms.¹² The idea is captured by an analogy offered by Matthen and Ariew (2009). Suppose two particles, p_1 and p_2 , in a container are moving away from one another at a constant velocity, such that their centre of mass, c , remains at a constant location with respect to the container. Now imagine that one particle, p_1 , contacts the wall of the container and rebounds back, while p_2 continues on its original trajectory. Instantaneously, c changes its location. Its change in location is a consequence of the change in relative motion of p_1 and p_2 . Yet there is no signal sent from p_1 and p_2 that causes the location of c to change. There couldn't be; since the change is instantaneous, a signal would have to travel faster than the speed of light. Nor does there need to be. The location of c is *entailed* by the locations of p_1 and p_2 . It is an analytic consequence: a function of the straight-line distance between p_1 and p_2 and their masses. Nothing is required to push, pull, attract or repel c directly, other than whatever causes the motions of p_1 and p_2 . In this respect, the motion of c is also a higher order effect of the motions of p_1 and p_2 . It is a change in the property of a system (the centre of mass of the system comprising p_1 and p_2) that results from causes acting on the system's components severally.

One of the striking features of higher order effects is that often large-scale trends or regularities can be seen by aggregating individual activities that could not be observed simply by following the behaviours of individuals separately. Consider Erwin Schrödinger's (1944) explanation of diffusion. If we place a drop of potassium permanganate in a volume of water it diffuses. The solution goes from a state in which there is a highly localized, deeply coloured purple area in an otherwise clear volume of water to an equilibrium state in which the entire solution is an evenly coloured light pink. Schrödinger explains diffusion in the following way. He asks us to imagine dropping a membrane into the solution some time before it has reached equilibrium. The motions of the individual permanganate molecules are random and unbiased. No molecule has more of a propensity to move in one direction than another. In such a situation, Schrödinger says, we would expect to observe that more molecules contact the barrier from the high density side than from the low density side. This is simply because there are more of them on the high density side. This means, then, that at any time prior to equilibrium, there will be more molecules traveling from high-density to low density than the other way around. The gross effect, summed over

¹¹Elliott Sober has helpfully pointed out that many causalists hold to some proper subset of these. That is true, but we maintain that only statisticalism is consistent with all four.

¹²The idea that this consequence is *analytic* is nicely presented by Neil Tennant (2014).

all these random motions, is that the solution will tend steadily, inexorably toward an equilibrium (even) distribution. Diffusion is a higher order effect. It is the consequence for an assemblage of the aggregate of the several causes of the motions of individual molecules.

Two morals can be drawn from the diffusion analogy, one metaphysical, the other methodological. The metaphysical lesson is that one need not posit higher-order causes in order to explain a higher order effect (Matthen 2010). There is no diffusive force. All the causing goes on at the individual level, despite the fact that the regularity is only observable at the ensemble level. The methodological moral is that higher order effects elicit a special kind of explanation: a higher order effect explanation. In order to explain the effect of the aggregate of individual-level causes we do not need to enumerate or articulate those individual-level causes severally. All that is required for the explanation is the *distribution* of these causes. Whatever the causes of molecular motion, there are more causing the motion from high density to low than *vice versa*.¹³

The statistical interpretation takes MS-selection to be a higher order effect. MS-selection is a change in the trait distribution of a population that results from the causes of individual birth, death, reproduction, mutation, and migration. Darwin's great conceptual breakthrough in the discovery of natural selection is that all the causes of descent with modification are already contained within the 'struggle for existence', the day-to-day activities of organisms. Nothing is required to cause evolution other than that which causes the survival, reproduction, inheritance, mutation, death, emigration and immigration of organisms in a population. Organisms survive, reproduce and die, enter and leave. As they do their trait tokens enter, stay in, or leave the population. As a consequence, some trait *types* increase in frequency relative to others.

To sum up, the statistical interpretation holds that the explanation of evolutionary population dynamics—that is to say, change in trait distribution—is a higher order effect explanation. Trait types change in their relative frequency as an analytic consequence of individual-level causes. A MS-model does not advert to the individual-level causes of population change. Instead it cites the *distribution* of those causes.

It is important to note that statisticalists do not deny that higher order effects can have individual-level causal explanations. For example, a satisfactorily complete individual-level causal explanation of why a particular *Drosophila melanogaster* population evolved a fuzzy thorax might cite a suite of multifarious causes for individuals and their kin over a number of generations. These may include the causes of survival, deaths, immigration, emigration, reproduction, and inheritance of characteristics, the typical effect of a fuzzy thorax on life history, etc. This would account for the change in lineage structure. It would be a perfectly good example of what Sober (above) refers to as an 'ecological' explanation. So long as we can, further, assign trait types to lineages, we can use this information to explain why change in lineage structure eventuates in change in trait distribution. In these individual-level causal explanations there is no need to refer to any probabilities or statistical distributions at all, there are no trait fitness ascriptions, or sample errors needed. What statisticalists deny is that such an explanation is an *MS-explanation*.

One might conclude that in the instances where we want to know why a particular population actually evolved a change in distribution of a trait (as in the *Drosophila* case) an individual-level explanation should be favored over an MS-explanation. (Rosenberg 2006, Ariew 1998, Sober 1993, Pence and Ramsey 2015). This might generally be correct. There may be no need, for example, for a MS-model to explain the evolution of industrial melanism in populations of *Biston betularia* in Birmingham (Kettlewell 1955). The observation that melanic individuals are less frequently preyed upon than 'typical' individuals suffices. But, this D-selection explanation is highly specific, and not generalizable. It does not tell us what would have happened if the specific history of births and deaths in Birmingham had been slightly different? Nor does it identify the commonalities between the change in this population and (for example) the increase in frequency of dark morph eastern grey squirrels (*Sciurus carolinensis*) in North America. This is where MS-explanations come in. They apply to all populations in which the growth rate of trait types vary, irrespective of how this variation is caused. After all, no two populations are alike in every causal detail,

¹³The distinction is presaged by Lewontin's (2001) distinction between 'variational' and 'transformational' explanations in biology.

but a wide range of them undergo the same type of changes in distribution patterns (Woodward 2003, Sober 1983, Ariew, Rice and Rohwer 2015). MS-models demonstrate that change in trait structure is sensitive to differences in the statistical distribution of a number of parameters.

Nor do statisticalists deny that higher order effects can sometimes have ensemble-level causes.¹⁴ For example, the bursting of a balloon under pressure is plausibly a higher order effect. It is a consequence of the motions of gas molecules with the membrane of the balloon. The balloon's burst is also caused by the pressure of the air, which is the cumulative effect of the motions of the molecules. Pressure is an ensemble level property that causes a higher order effect.¹⁵ The difference between pressure bursting a balloon and MS-selection is that the former is not an *analytic* consequence of individual-level processes. The aggregate motions of individual molecules do not *entail* the bursting of the balloon; they cause it. Pressure is the *sum* of forces exerted by individual molecules (per unit area, per unit time). The balloon pops when the total force applied to it exceeds a certain threshold value.¹⁶

Trait fitness is a primitive.

The crucial explanatory concept in MS-models is trait fitness. Trait fitness estimates the relative growth rate of a trait type.¹⁷ The statistical interpretation holds that trait fitness is a primitive concept. It has no definition expressed in terms of any other parameter. In particular, (unlike other aggregate-level parameters, e.g. pressure and temperature) it has no definition in terms of the causal properties of individuals that make up an ensemble. In contrast, most causalists seem to assume—and sometimes explicitly aver—that the concept of trait fitness is derivative. It is defined as a function of individual causal properties (Beatty 1998, Beatty and Finsen 1998, Pence and Ramsey 2015). It is frequently said that trait fitness supervenes on vernacular fitness (Rosenberg 2008).

Philosophical folklore has it that the fitness of a trait is the mean of the vernacular fitnesses of those individuals that possess the trait. Vernacular fitness in turn is thought of as a propensity of an individual (usually the propensity to reproduce). As vernacular fitness is a causal property, the supposition goes, the average of these properties must be a causal property too. It might be thought of as the general contribution that a trait token makes to the fitness of an individual that bears it (Gillespie 1977). Perhaps this is analogous to the average effect that, say, a container of steel pipes makes to the displacement of any ship that it might be loaded onto, or the average contribution that a litre of petrol makes to the distance a car of a certain model can travel. This is an appealing line of thought, but it should be resisted for two reasons. The first is that trait fitness is typically not equivalent to average vernacular fitness. The second concerns the role that trait fitness plays in MS-models.

Typically, average individual reproductive output is too crude as a representation of the tendency of a trait type to increase or decrease in relative frequency (Gillespie 1977, Lande 1978). Demographic factors, such as population size, age structure, and whether the population is growing or diminishing, also affect the rate of change in frequency of a trait type. For instance, for organisms with continuous reproduction and overlapping generations the average age at which members of a type produce offspring makes a difference to changes in population composition.

So even though one trait might be changes in population that individuals with the trait tend to survive or reproduce better than the alternative, the overall expectation as to which trait will increase in the population must take into account all biasing effects from all processes causing evolutionary change. (Earnshaw-Whyte 2012: 409)

¹⁴See Shapiro and Sober (2007) and Millstein (2006).

¹⁵Jackson and Pettit (1990) think that explanations that advert to pressure are mere 'program explanations', placeholders for the real, individual-level causal explanations. We're happy to acknowledge that these are perfectly good causal explanations.

¹⁶For more on this see Hacking (1990), and Ariew, Rice and Rohwer (2015) on statistically autonomous explanation.

¹⁷Alternatively, the rate of change in mean character value of a trait type

For another example, variance in reproductive output can make a difference to the change in relative frequency of two trait types even if the types have the same mean reproductive output. Furthermore, overall population size affects the degree to which variance affects changes in trait frequency (Sober 2000, Orr 2007, Walsh 2015). As Sober (2000) astutely points out, if the rate of growth of a trait type depends upon the number of organisms in the population (whether they possess the trait or not), and average vernacular fitness doesn't, then the growth rate of a type is not a function of the average fitnesses of individuals possessing tokens of that type.

These considerations and others (discussed in Ariew and Lewontin, 2004), suggest that the rate of growth of a trait type is dependent upon all manner of causal, demographic, and statistical factors that are extraneous to the average reproductive output of the individuals possessing a given trait.

Earnshaw-Whyte calls the average vernacular fitness of those organisms with a given trait 'narrow trait fitness' in contradistinction to 'broad trait fitness', which is the 'all-things-considered' tendency of a trait type to increase or decrease in its relative frequency. He argues that the 'narrow' fitness of a trait is typically insufficient to predict or explain the change in relative frequency of that type.

By itself, this does not establish that the distribution of trait fitnesses cannot be defined in terms of the properties of individuals in the population. Demographic factors, like population growth, population size, age structure, reproductive schedule, and migration rate may well contribute to trait fitness. But these phenomena are fixed by the properties of individual organisms. That being so, the distribution of trait fitnesses might still supervene on the properties of individuals in a population, even if it doesn't *oesnn if ie* on vernacular fitnesses. In that event, trait fitness distribution would not be a primitive. It would have a definition given exhaustively in terms of the properties of individual organisms.

But a consideration of the role that trait fitness distribution plays in MS-models shows this idea to be implausible (Ariew and Ernst 2009). The processes that cause populations to change are multifarious, spatially and temporally variable, and stochastic. MS-models only ever encompass a simplified, idealized, probabilistic summation of the suite of causes that a population might encounter (Levins 1966, Weisberg 2006). Consequently, the distribution of trait fitnesses in a model renders a probability distribution over the range of possible population changes. We can think of an actual population—the population whose dynamics is being modeled—as a sample drawn from an indefinitely large source of possible populations, each of which is described by the same model. Now, one obvious fact about samples is that they may misrepresent their source population. Two samples may be identical, and yet be drawn from source populations with different statistical properties. It is entirely plausible, for example, that two runs of ten flips of different coins may both produce five heads and five tails, and yet one coin be fair and the other biased. Likewise, two populations could have an identical arrangement of individual-level causal properties and yet differ in their distribution of trait fitnesses. That being so, trait fitness does not supervene on the sum of actual individual-level causes in a population.

Two things follow. The first is that the suite of actual individual-level causes of organisms leaving, staying in, or entering a population underdetermines the distribution of trait fitnesses. Trait fitness thus cannot be defined in terms of the properties of actual individuals in a population. The second is that trait fitness is an abstraction. It is a statistical property of an *idealized* population and cannot be defined in terms of the properties of actual individuals in a concrete population. Trait fitness is thus best considered to be a primitive theoretical concept.

This claim should be differentiated from two previous attempts to establish the primitiveness of trait fitness (See Sober (2013), and Rosenberg and Williams (1986)). Elliott Sober (2013) articulates a position in which trait fitnesses play a primitive role in population models. This sounds pretty close to the statisticalist view. However, Sober diverges from statisticalism in three crucial ways. He holds that: (i) trait fitnesses can be *defined* in terms of vernacular fitness (ii) vernacular fitness has no explanatory uses, and (iii) *variation* in trait fitness causes evolution. Statisticalists deny all three of these theses. Similarly, Rosenberg and Williams (1986) posit the primitiveness of trait fitness, in order to include it as a term in an axiomatic formal syntactic account of the logical structure of evolutionary theory. This is no part of the statisticalist strategy. The statisticalist position on trait fitness derives from an observation of their role in MS-models.

Population dynamics models are substrate neutral.

A model is a partial, indirect representation of an object of study, in this case, a biological population. Presumably, no single workable model could include all the highly specific causal details of change in population structure *and* capture the dynamics of populations in their full generality (even if it were desirable to do so). So, we must make choices about what factors to represent, which details to idealize and which to leave out. These choices are guided by pragmatic considerations. There are a number of strategies available to modelers (Levins 1966). One is to sacrifice realism for generality and precision (Weisberg 2006). This is what MS-models do. They render highly general and precise descriptions of the dynamics of populations—construed as ensembles of abstract types—at the cost of leaving out all the strictly biological details.

MS-models make only minimal metaphysical demands on the populations to which they are applied. The conditions for applying a MS-model may be realized in biological populations, but virtually nothing of what makes a biological population *biological* is explicitly represented in the model (Stegenga 2014). Matthen and Ariew (2002), for example, argue that one can think of a MS-model as an instance of Li's theorem:

Li's theorem: in a subdivided population the rate of change in the overall growth rate is proportional to the variance in growth rates.

Li's theorem identifies a regularity, but not an expressly *biological* one.

Winther et al (2015) draw the helpful distinction between biological populations and theoretical populations. Biological populations may all be theoretical populations, but not *vice versa*.¹⁸ Biological populations are those in which organisms partake in the processes of birth, death, and reproduction (*inter alia*). Theoretical populations are those that meet the minimal conditions required for the application of a MS-model. These conditions are so minimal that MS-models are equally applicable to a wide range of non-biological ensembles. As Li's theorem suggests, they can apply to any system in which persistent types vary in their growth rates. Darwinian selection in biological populations is one way of meeting the conditions, but there are many others. Orr (2007) and Walsh (2015), for example, discuss how the models that describe the change in relative frequency of traits in a population undergoing selection can also be used to represent the change in relative growth rates of competing investment portfolios. Hodgson and Knudsen (2010) make the point that evolutionary models apply to business practices, legal systems and technological change, just as well as to biological populations. In this sense, MS-models are substrate neutral (Matthen and Ariew 2002). On account of their substrate neutrality, MS-explanations say nothing about the distinctly biological processes that bring about population change.

This does not make MS-models useless, far from it. The loss of specific causal detail is more than compensated by a gain in generality. MS-models enable biologists to identify, measure, and project regularities at the level of the dynamics of ensembles that cannot be described without ascending to an ontology of abstract trait types and their growth rates (Walsh, Lewens, Ariew 2002, Walsh 2003). They articulate the ways in which the rate, direction and magnitude of change in the trait distribution of an assemblage counterfactually depend upon its statistical structure. These regularities are largely insensitive to the specifically biological details. MS-models capture evolutionary phenomena in their full generality, as they apply equally to populations of *Latimeria* and *Laminaria*, bacteria colonies and Bactrian camels, and much more besides.

MS-Selection and drift are model-relative.

As we have said, a MS-model is a highly simplified, abstract and idealized representation of biological reality. The choices that an investigator makes in constructing a model are guided by her objectives. As objectives may vary, a population may have many different MS-models depending on which features the

¹⁸Millstein (2006) and Gildenhuys (2014), for example, erroneously suppose that being a biological population is a necessary condition for the application of an MS-model.

Table 1: Change in frequency of traits H and T over a run of eight years

Trait type	# of offspring		Relative frequency at the end of year ...							
	warm	cool	1	2	3	4	5	6	7	8
H	6	4	.60	.69	.77	.69	.77	.84	.88	.84
T	4	5	.40	.31	.23	.31	.23	.16	.12	.16

investigator chooses to accentuate, ignore, or idealize. Trait fitnesses are assigned relative to a MS-model. If MS-selection and drift are defined in terms of variation in trait fitnesses, then they too are attributed to a population relative to a model.

For certain purposes it might be prudent to describe a population as varying in trait fitnesses, while for others the most apt description of the same population may be one in which there is no variation in trait fitness.¹⁹ We can only say, then, that a population varies in its trait fitness *with respect to some model* or other. Consequently, we can only say that a population is undergoing MS-selection, or drift, or both, *relative to a MS-model*. It follows that there is no model-independent fact of the matter whether a population is undergoing MS-selection, drift, or both. A toy scenario might help to illustrate.²⁰

Suppose we are following the growth of a founding population of 20 asexually reproducing individuals in which there are two true-breeding variant traits $H:T$. Reproduction is synchronized, and each individual reproduces once and then dies, so there are no overlapping generations. There are 10 H individuals, and 10 T , each produces on average 5 offspring per year. The types are differentially sensitive to temperature. H individuals are fitter than T individuals in warmer than average years (where H individuals produce 6 offspring and T individuals produce 4). T individuals are fitter than H individuals in cooler than average years (where H individuals produce 4 offspring and T individuals produce 6). During our eight year study there is a run of six warm years (w) and two cool (c) as follows: w, w, w, c, w, w, c, w . The change in trait distribution is given in Table 1.

The question is how we should describe the dynamics of this population. On one hand, this looks to be a genuine case of MS-selection. The fitnesses of the trait types vary over the run of years: H is fitter than T (average reproductive output: $H=5.5$; $T=4.5$). Variation in fitness predicts and explains the increase in H relative to T . On the other hand, this also looks like a case of drift without selection. *Ex hypothesi*, over long stretches of time the fitnesses of H and T are equal. We should predict no net increase in the frequency of one over the other. However, this sample of six warm years and two cool years misrepresents the long-term run of years. There is sample error here—drift.

So, there are two distinct MS-models of this population. There is a ‘selection’ model in which $wH > wT$, and a ‘drift’ model in which $wH = wT$. The important point is that each of these models is well-suited to a different explanatory project. The description that attributes MS-selection and no drift to the population tells us that the change in population structure is entirely predictable. It further demonstrates what this population, over this stretch of time, has in common with other populations in which trait fitnesses vary in the same way. Yet, if we were to project this variation in fitness over a prolonged period of years, this model would *misrepresent* the change in the population. The description in which there is drift and no MS-selection, by contrast, underscores the fact that this run of eight years is something of a statistical anomaly. In any arbitrarily chosen run of eight years it would be unlikely that we should find such a skew in the conditions beneficial to H and detrimental to T . The no-fitness-variation description demonstrates what this run of eight years has in common with any other arbitrarily chosen sequence of years. While the $wH = wT$ fails to predict the population change over the run of eight years, it should fare better over other arbitrarily chosen runs of eight years, over extended periods of time, and for other populations of H and T individuals. So, both models are correct. The upshot is that a population can only be said to be undergoing MS-selection or drift *relative to a model*.

¹⁹A number of arguments have been offered by Walsh (2007, 2010, 2013, 2015).

²⁰Adapted from Walsh (2013)

What is *not* model-relative, is the change in trait structure of the population. Any way you slice the cake, there are more *H*s than *T*s, at the end of the eight years, and *that* requires an explanation. Both causal and statistical interpretations agree that MS-models explain this. The statistical interpretation further holds that whether MS-selection is occurring in a population simply depends on how the population is described.

The model relativity of MS-selection and drift marks a watershed between the statistical and causal interpretations. The model-relativity of MS-selection and drift sits uneasily with the causal interpretation, while it is perfectly consonant with statisticalism.

The tension between model-relativity and the causal interpretation arises from the commonly-held intuition that causal relations are objective, description-independent features of the world. That is to say for any *X* for any *Y*, there is an objective matter of fact whether *X* is a cause of *Y*, irrespective of how *X* and *Y* are described.²¹ That being so, if MS-selection and drift are causes, then for any population change, there should be a description-independent fact of the matter whether it is a case of selection, or drift, or a combination of the two. If there is a model-independent matter of fact in our toy scenario that $wH > wT$, for example, then it is just not true that $wH = wT$. The models cannot both apply to the same population consistently. The causal interpretation has two options. It might accept that both models describe this population, and thereby accept a contradiction—that selection objectively both is and isn't the cause of the population change. Alternatively, it might (arbitrarily) reject one or the other models as false (see Abrams 2013), and thereby forfeit the explanation that it offers (Walsh 2007).²²

The statistical interpretation encounters no such difficulty. It denies that MS-selection and drift are causes of population change. They are merely features of a statistical description of the population. Whether a population is undergoing MS-selection, or drift, or both is only a fact about how the population is described. One and the same system may be susceptible to different statistical descriptions. As statistical descriptions, both the $wH > wT$, and $wH = wT$ models are correct. Their correctness is borne out in their empirical success. According to one description, there is MS-selection and no drift; according to the other, there is drift and no MS-selection. But there is no inconsistency, as a population can be said to be undergoing MS-selection, drift, or both *only relative to a model*.

Description dependence underscores a significant contrast between Darwinian natural selection and Modern Synthesis natural selection. D-selection, just like change in trait structure, is description *in*-dependent. D-selection occurs when organisms survive and reproduce differentially on account of their vernacular fitnesses (which as we have seen are causal propensities). Vernacular fitnesses are *non*-model-relative. If two D-models assign incompatible vernacular fitnesses to the organisms in our toy scenario, for example, then at least one of them is wrong. The error will show up in the model's lack of empirical success.²³ This is a crucial difference between D-selection and MS-selection that usually goes unnoticed. Part of the confusion surrounding the causalist/statisticalist debate, we conjecture, stems from a general failure to distinguish model-relative MS-selection (and drift) from description-independent D-selection, and a concomitant failure to draw the distinction between MS-models that ascribe trait fitnesses and D-models that ascribe vernacular fitnesses.

²¹Caution is required here. Many approaches to modeling causation, e.g. Woodward (2003) allow that different causal models may describe the same system differently. For example, one model may represent electromagnetic force and gravitational force, but not displacement, mass or charge as distinct causes, while another may represent displacement, mass and charge as causes, without explicitly representing gravitational or electromagnetic force. But this does not entail that these causal relations are model-relative. It simply means that the objective causal relations can be described in multiple, mutually consistent ways. We thank an anonymous referee for raising this issue.

²²A third, we suggest heroic, alternative would be to reject the description independence of causal relations. See Northcott (2010).

²³The error, of course, is not statistical error (Walsh, Lewens, and Ariew 2002). Drift does not figure in D-models.

4 Conclusion

The statistical interpretation holds that MS-models of evolutionary change are statistical models. They cite trait fitnesses, the differential growth rates of abstract trait types. Trait fitness is not a causal property of a concrete entity. It is a statistical property of an abstract entity. MS-Selection occurs when there is variation in trait fitnesses in a population. MS-selection is not to be conflated with the Darwinian process of selection (D-selection), in which individuals survive and reproduce differentially according to their own causal propensities. D-selection may be *realized* as MS-selection (and *vice versa*), but they are different processes. Much of the debate concerning the statistical interpretation of evolutionary theory over that last 15 years stems from a tendency to conflate these two processes. Moreover, the debate has systematically conflated two kinds of evolutionary models, D-models and MS-models. The distinctions are important. D-selection models are causal; MS-selection models are not. MS-selection is model relative; D-selection is not.

The statistical interpretation is exclusively a claim about *MS-selection and its models*. It has nothing to say about D-selection, or D-models. It makes the following claims:

Positive Claim #1: MS-models cite statistical properties of trait types.

Positive Claim #2: MS-models explain changes in trait distribution.

Negative Claim: MS-models do not cite the causes of population change.

These claims are supported by the following commitments:

1. Natural Selection is a higher order effect.
2. Trait fitness is primitive.
3. MS-models are substrate neutral.
4. MS-selection and drift are model-relative.

Acknowledgments

This paper began as a response to a request from Alex Rosenberg to set out the precepts of statisticalism. We thank Alex for providing the impetus. We benefited from helpful feedback from audiences and discussion groups at ISHPSSB Montréal, University of Toronto, Duke University, Washington University, University of Wisconsin-Madison. In particular we would like to thank Elliott Sober, Fermín Fulda, Cory Lewis, Eugene Earnshawe-Whyte, and Richard Lauer for helpful feedback.

Literature cited

- Abrams, Marshall. 2013. "Populations and Pigeons: Prosaic pluralism about evolutionary causes" *Studies in History and Philosophy of Biological and Biomedical Sciences* 44:294–301.
- Ariew, André. 1998. "Are Probabilities Necessary for Evolutionary Explanations?" *Biology and Philosophy* 13(2): 245–53.
- Ariew, André. 2003. "Ernst Mayr's 'Ultimate/Proximate' Distinction Reconsidered and Reconstructed." *Biology and Philosophy* 18(4): 553–65.
- Ariew, André. 2008 "Population Thinking." *Oxford Handbook of Philosophy of Biology*. Oxford: Oxford University Press. 64–86.
- Ariew, André and Zachary Ernst. 2009. "What Fitness Can't Be." *Erkenntnis* 71(3): 289–301.

- Ariew, André and R. C. Lewontin. 2004. "The Confusions of Fitness." *The British Journal for the Philosophy of Science* 55(2): 347–63.
- Ariew, André, Collin Rice, and Yasha Rohwer. 2015. "Autonomous-Statistical Explanations and Natural Selection." *The British Journal for the Philosophy of Science* 66(3): 635–58.
- Beatty, John. 1998. "Fitness: Theoretical contexts." In E. F. Keller & E. A. Lloyd (Eds.), *Keywords in Evolutionary Biology*. Cambridge, MA: Harvard University Press.
- Beatty, John and S. Finsen. 1998. "Rethinking the Propensity Interpretation: A Peek Inside Pandora's Box." In *What the Philosophy of Biology Is: Essays Dedicated to David Hull*, edited by Michael Ruse. Dordrecht: Kluwer Academic Publishers. 17–30.
- Brandon, Robert N. 2006. "The Principle of Drift: Biology's First Law." *The Journal of Philosophy* 103(7): 319–35.
- Darwin, Charles. 1859. *The Origin of Species By Means of Natural Selection*.
- Earnshaw-Whyte, Eugene. 2012. "Increasingly Radical Claims About Heredity and Fitness." *Philosophy of Science* 79(3): 396–412.
- Endler, John. 1986. *Natural Selection in the Wild*. Monographs in Population Biology 21. Princeton: Princeton University Press.
-  Fisher, Ronald Aylmer. 2000. *The Genetical Theory of Natural Selection: A Complete Variorum Edition*. Oxford University Press.
- Gildenhuys, Peter. "Arbitrariness and Causation in Classical Population Genetics." *British Journal for the Philosophy of Science* (2014) 65: 429–444.
- Gillespie, John H. 1977. "Natural Selection for Variances in Offspring Numbers: A New Evolutionary Principle." *The American Naturalist* 111(981): 1010–1014.
- Grene, Marjorie. 1961. "Statistics and Selection." *British Journal for the Philosophy of Science* 45: 25–42.
- Hacking, Ian. 1990. *The Taming of Chance*. Cambridge University Press.
- Hitchcock, Christopher and Joel D. Velasco. 2014. "Evolutionary and Newtonian Forces." *Ergo* 1.
- Hodgson, Geoffrey, M. and Thorbjørn Knudsen. 2010. *Darwin's Conjecture: The Search For General Principles Of Social And Economic Evolution*. Chicago: Chicago University Press.
- Jackson, Frank and Philip Pettit. "Program Explanation: A General Perspective." *Analysis* (1990): 107–117.
- Kettlewell, H. B. D. 1955. "Selection experiments on industrial melanism in the *Lepidoptera*." *Heredity* 9: 323–342.
- Lande, Russell. 1976. "Natural Selection and Random Genetic Drift in Phenotypic Evolution." *Evolution* 30(2): 314–334.
- Lennox, James G. 1991. "Darwinian Thought Experiments: A Function for Just-So Stories." In *Thought Experiments in Science and Philosophy*, edited by Tamara Horowitz and Gerald J. Massey. 223–45.
- Levins, Richard. 1984. "The Strategy of Model Building in Population Biology." (1966) Reprinted in: Sober E (ed) *Conceptual Issues in Evolutionary Biology*, 1st ed. Cambridge MA: MIT Press. 18–27
- Levins, Richard and Richard C Lewontin. 1985. *The Dialectical Biologist*. Harvard University Press.
- Lewontin, Richard C. 2001. *The Triple Helix: Gene, Organism, and Environment*. Harvard University Press.

- Matthen, Mohan. 2009. "Drift and Statistically Abstractive Explanation." *Philosophy of Science* 76(4): 464–87.
- Matthen, Mohan. 2010. "What is Drift? A response to Millstein, Skipper, and Dietrich." *Philosophy & Theory in Biology* 2(e102): 1–6.
- Matthen, Mohan and André Ariew. 2002. "Two Ways of Thinking About Fitness and Natural Selection." *The Journal of Philosophy* 99(2): 55–83.
- Matthen, Mohan and André Ariew. 2005. "How to Understand Casual Relations in Natural Selection: Reply to Rosenberg and Bouchard." *Biology and Philosophy* 20(2): 355–64.
- Matthen, Mohan and André Ariew. 2009. "Selection and Causation." *Philosophy of Science* 76(2): 201–24.
- Millstein, Roberta L. 2002. "Are Random Drift and Natural Selection Conceptually Distinct?" *Biology and Philosophy* 17: 33–53.
- Millstein, Roberta L. 2006. "Natural Selection as a Population-Level Causal Process." *The British Journal for the Philosophy of Science* 57(4): 627–53.
- Morrison, Margaret. 2002. "Modelling Populations: Pearson and Fisher on Mendelism and Biometry." *The British Journal for the Philosophy of Science* 53(1): 39–68.
- Northcott, Robert. 2010. "Walsh on Causes and Evolution." *Philosophy of Science* 77: 457–467.
- Nowak, Martin A. 2006. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press.
- Okasha, Samir. 2015. "Population Genetics." *Stanford Encyclopedia of Philosophy*. Fall 2015 Edition, edited by Edward N. Zalta.
<http://plato.stanford.edu/archives/fall2015/entries/population-genetics>
- Orr, H. Allen. 2007. "Absolute Fitness, Relative Fitness, and Utility." *Evolution* 61: 2997–3000.
- Pence, C. H. and G. Ramsey. 2015. "Is Organic Fitness at the Basis of Evolutionary Fitness?" *Philosophy of Science* 82: 1081–1091.
- Rosenberg, Alexander. 2008. *Darwinian Reductionism: Or, How to Stop Worrying and Love Molecular Biology*. University of Chicago Press.
- Rosenberg, Alexander and Mary Williams. 1986. "Fitness as Primitive and Propensity." *Philosophy of Science* 53(3): 412–18.
- Schrodinger, Erwin. 1989. *Statistical Thermodynamics*. Dover Publications.
- Shapiro, Larry and Elliott Sober. 2007. "Epiphenomenalism—the Do's and Don'ts." In *Studies in Causality: Historical and Contemporary*, edited by G. Wolters and P. Machamer. Pittsburgh: University of Pittsburgh Press. 235–64.
- Sober, Elliott. 1983. "Equilibrium Explanation." *Philosophical Studies* 43: 201–10.
- Sober, Elliott. 1993. *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. University Of Chicago Press.
- Sober, Elliott. 2001. "The Two Faces of Fitness." In *Thinking About Evolution: Historical, Philosophical, and Political Perspectives*, edited by Rama S. Singh, Costas B. Krimbas, Diane B. Paul, and John Beatty. 309–21.
- Sober, Elliott. 2013. "Trait Fitness Is Not a Propensity, But Fitness Variation Is." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44(3): 336–41.

- Stegenga, Jacob. 2014. "Population Pluralism and Natural Selection." *The British Journal for the Philosophy of Science* 67(1): 1–29.
- Stephens, Christopher. 2004. "Selection, Drift, and the 'Forces' of Evolution." *Philosophy of Science* 71(4): 550–70.
- Tennant, Neil. 2014. "The Logical Structure of Evolutionary Explanation and Prediction: Darwinism's Fundamental Schema." *Biology and Philosophy* 29: 611–655.
- Walsh, Denis M. 2000. "Chasing Shadows: Natural Selection and Adaptation." *Studies in History and Philosophy of Biological and Biomedical Sciences* 31(1): 135–53.
- Walsh, Denis M. 2004. "Bookkeeping or Metaphysics? The Units of Selection Debate." *Synthese* 138(3): 337–61.
- Walsh, Denis M. 2007. "The Pomp of Superfluous Causes: The Interpretation of Evolutionary Theory." *Philosophy of Science* 74(3): 281–303.
- Walsh, Denis M. 2010. "Not a Sure Thing: Fitness, Probability, and Causation." *Philosophy of Science* 77(2): 147–71.
- Walsh, Denis M. 2010. "Two Neo-Darwinisms." *History and Philosophy of the Life Sciences* 32(2/3): 317–39.
- Walsh, Denis M. 2012. "The Struggle for Life and the Conditions of Existence: Two Interpretations of Darwinian Evolution." In *Evolution 2.0: Implications of Darwinism in Philosophy and the Social and Natural Sciences*, edited by Martin H. Brinkworth and Friedel Weinert. Springer. 191–209.
- Walsh, Denis M. 2013. "Descriptions and Models: Some Responses to Abrams." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44(3): 302–8.
- Walsh, Denis M. 2014. "Variance, Invariance and Statistical Explanation." *Erkenntnis* 80(Suppl. 3): 1–21.
- Walsh, Denis M. 2015. "Variance, Invariance, and Statistical Explanation." *Erkenntnis* 80: 469–489.
- Walsh, Denis M, Tim Lewens, and André Ariew. 2002. "The Trials of Life: Natural Selection and Random Drift." *Philosophy of Science* 69(3): 429–46.
- Weisberg, Michael. 2006. "Forty Years of 'The Strategy': Levins on Model Building and Idealization." *Biology and Philosophy* 21: 623–645.
- Winther, Rasmus Grønfeldt, Ryan Giordano, Michael D. Edge, and Rasmus Nielsen. 2015. "The Mind, the Lab, and the Field: Three Kinds of Populations in Scientific Practice." *Studies in History and Philosophy of Biological and Biomedical Sciences* 52: 12–21.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.

© 2017 Author(s)

This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International license (CC BY-NC-ND 4.0), which permits anyone to download, copy, distribute, or display the full text without asking for permission, provided that the creator(s) are given full credit, no derivative works are created, and the work is not used for commercial purposes.