

Pressing Christie, Brusse, et al.'s Objection: Why Single Out Selected Effects?

Aliya R. Dewey

Department of Philosophy, University of Arizona

Christie et al. (2022) argue that selected effects are *insufficient* to explain the prevalence of traits when selection is heterogenous. One could object that it's useful to ground functions in selected effects so long as selected effects are *necessary* to explain the prevalence of traits. This raises a challenging question: what justifies singling out selected effects from other factors that are *necessary* to explain the prevalence of traits when selection is heterogenous? I consider three answers: selected effects are the only factors that explain the prevalence of traits and (a) can plausibly ground functions, (b) enable biological generalization, and (c) enable representational explanation. I argue that none of these answers is satisfactory. Only the third answer might be able to justify singling out any selected effects at all (when selection is heterogenous): the small share of selected effects that ground representations of causal particulars. Thus, I agree with Christie et al. that heterogenous selection is a serious problem for selected effects theories of function.

Keywords: selected effects; functions; generalization; representation; naturalism

§1. Selected effects theories of function

Other commentators have argued that Christie et al.'s (2022) argument mistakes the explanatory aims of selected effects theories (SETs) of function (Garson, 2022; Kingsbury, 2022; Okasha, 2022; Shea, 2022).¹ I agree with this concern, but I think that Christie et al.'s argument survives with some friendly reconstruction, so I'll start by reviewing what I take to be the (direct) explanatory aim of SETs.

SETs are meant to explain which natural properties ground functional properties. Suppose that the function of a trait T is to A. Philosophers of biology ask: what could make that fact true? SETs explain that this fact is grounded in a fact about selection: the function of any trait T is to A iff (and because) T was selected for A-ing. In other words, A-ing is T's function because A-ing is T's selected effect (SE). For example, the function of zebra stripes might be to deter biting flies iff (and because) the zebra stripes were selected for deterring biting flies.

By comparison, SETs are *not* meant to explain the prevalence of traits in a population (despite how Christie et al. sometimes suggest). After all, they would do this poorly. Suppose that a trait T is present in a population. What makes that true? One plausible explanation is that T was selected for A-ing. Whether T being selected for A-ing grounds the further fact that T has the function to

¹ All citations refer to other works in this volume, unless otherwise noted.

A is irrelevant to that question. After all, suppose that SETs are false (maybe the organizational theory of function is true instead): the fact that T has the function to A isn't grounded in the fact that T was selected for A-ing. This would make no difference to our ability to explain the prevalence of T in a population: it would still be due to the fact that T was selected for A-ing.

Of course, biologists often do use 'T has the function to A' as a shorthand for saying that 'T was selected for A-ing'. And SETs are supposed to vindicate this practice: they explain the fact that T has the function to A is grounded in the fact that T was selected for A-ing. This is often touted to be the crowning achievement of SETs (Garson, 2022). But the consequences wouldn't be dire if SETs were false: this shorthand would be inaccurate but probably harmless. For this reason, I'm much less impressed with the naturalization project than Garson is (more on that in §5).

§2. Why single out selected effects?

A critical part of Christie et al.'s argument (not all of it) is their claim that SETs don't *completely* explain the prevalence of traits in populations when the selection environment is heterogeneous. If there's an arms race between zebra stripes and the preference traits of biting flies, the possible fact that zebra stripes were selected for deterring biting flies is only part of the explanation. The rest of an adequate explanation should mention that the prevalence of zebra stripes would select for preferences for striped surfaces in biting flies. Then lacking zebra stripes would be selected for deterring biting flies. An adequate explanation of the prevalence of zebra stripes in a population should include both the effects that it's selected *for* causing (which ground its SE function) and the effects that it's selected *against* causing (which don't ground its SE function).

Other commentators have argued that Christie et al.'s argument is a non-sequitur because SETs were never meant to completely explain the prevalence of traits (Garson, 2022; Kingsbury, 2022). There's more to Christie et al.'s argument, though. They recognize that SETs single out SEs as grounds for functions. This raises a question for SETs: what justifies singling out SEs from the other things that explain the prevalence of a trait in a population? Why single out the effect of deterring biting flies, e.g., when the opposite effect of attracting biting flies is equally important to explaining the prevalence of zebra stripes in a population? The heterogeneity of selection is indispensable to this challenge. When selection is homogenous, it's straightforward to justify singling out SEs from other things: they are sufficient to explain the prevalence of a trait in a population. When selection is heterogenous, though, it isn't straightforward to justify singling out SEs from other things: SEs aren't any more relevant to explanations of trait prevalence than some non-SE things.

§3. Because they ground functions

One response to Christie et al.'s challenge is to say that we are justified in singling out SEs from other causal factors because SEs are the only causal factors that could ground any recognizable notion of function. Stripes may have the positive effect (on their prevalence) of deterring biting flies sometimes and the negative effect (on their prevalence) of attracting biting flies other times, but only the positive effect of deterring biting flies is a plausible ground for functions. If we're interested in grounding functions in the factors that explain trait prevalence, then only one of those effects will do: deterring biting flies. That justifies singling it out. Kingsbury (2022) puts the point

most clearly: “although stripes have not consistently contributed to fitness, the SET is only interested in what they have done when they *have* contributed to fitness.”

However, this response undermines the motivation for SETs. After all, SETs originally offered a justification for singling out functions: SEs warrant singling out and SEs ground functions. This could still be true when selection is homogenous.² But Christie et al. have argued that selection is often heterogenous. Explanations of trait prevalence don’t justify singling out SEs in such cases. If SETs are right that functions are grounded in SEs, explanations of trait prevalence won’t justify singling out functions either. This isn’t a fatal problem for SETs: they can find other ways to justify singling out SEs and functions. But “until this work is done [SETs run] the risk of merely hand-waving at natural selection to lend an air of respectability to normative intuitions” (Christie et al., 2022).

§4. Because they enable biological generalization

Another response to Christie et al.’s challenge, due to Neander (2017), is to say that we are justified in singling out SEs and SE functions because doing so is necessary for solving the “generalization problem” in biology. According to Neander, the generalization problem is that the huge number and diversity of possible activities for biological systems outstrips our ability to explain all, or even a representative sample, of them. For example, there is an arbitrary number of possible things that zebra stripes *could* do, from obscuring their boundaries to creating uneven heat gradients along their bodies. We lack the resources to build a general theory of zebra stripes by explaining all, or even a representative sample, of all possible things they could do.

Neander argues that this generalization problem is limited to dysfunctional activities: functional activities are rare and uniform. This is the Anna Karenina principle: “All happy families are alike; each unhappy family is unhappy in its own way” (Tolstoy, 1877). Why do biological systems satisfy this principle? She argues that they satisfy the first part of the principle because “selection has pushed many adaptive traits to fixation and so towards something that approximates a single species design” (1158). Next, she argues that they satisfy the second part of the principle because “there is no limit on pathological variation other than the laws of nature and death, whereas normal functioning is more highly constrained” (1159).

If this is true, we can solve the generalization problem by (a) singling out the small, uniform set of SE functional activities for explanation and (b) abstracting away from the huge, diverse set of its SE dysfunctional activities. But this solution only works if SE functional activities are *relevant* to a general theory of function-bearing systems and SE dysfunctional activities are not. That is, it only works if SE functional activities warrant singling out from SE dysfunctional activities in the first place. This is where Christie et al. can press their challenge: if functions are SEs, they don’t warrant singling out from dysfunctional activities unless selection is homogenous. So, Neander’s solution to the generalization problem presupposes that SEs and SE functions already warrant singling out, which, Christie et al. effectively argue, presupposes the typically false assumption that selection is homogenous.

² I reject this move for a different worry: if all functions are SE functions and SE functions are just SEs, then functions are ontologically redundant vis-à-vis SEs. If this is true, then we aren’t justified in singling out functions after we’ve singled out SEs. Hence, Garson’s monist type of SET strikes me as a pyrrhic victory.

§5. Because they enable representational explanation

A final response to Christie et al.’s challenge, which may be the most popular, is to say that we are justified in singling out SEs and SE functions because doing so is necessary for representational explanation of behaviour (Garson, 2022; Kingsbury, 2022; Okasha, 2022; Shea, 2022). After all, adequate explanations of behaviour often require us to attribute *representations* to cognition and representations are often grounded in functions, especially SE functions. For example, teleosemantic theories roughly claim that a cognitive state S represents a content C iff (and because) the tokening of S has the function to stand in some causal or correlative relation with the tokening of C. This claim also allows us to distinguish accurate representations from inaccurate ones: any representation is accurate iff (and because) it achieves the SE function that grounds it.

Suppose that teleosemantic theories are necessary for explaining why adequate explanations of behaviour often require us to attribute representations to agents, regardless of whether selection is homogenous or heterogenous (Shea).³ Then we’d be justified in singling out SEs and SE functions that ground representations—even when selection is heterogeneous. If this all works out, it would be a win for SETs: they could justify singling out SEs that ground SE functions that, in turn, ground representations. But it would be a very small win: most SEs and SE functions (e.g., the putative function of zebra stripes to deter biting flies) exist in heterogenous environments yet have nothing to do with representations. So, representational explanations couldn’t possibly justify singling out the majority of SEs or SE functions.

But matters are even worse than this for SETs. Teleosemantic accounts are suitable for explaining cognitive representations of *causal particulars*, which cognitive scientists tend to care about. After all, the tokening of cognitive representations can stand in causal or correlative relations with the tokening of causal particulars. But teleosemantic accounts aren’t suitable for explaining cognitive representations of non-causal particulars and universals, which philosophers of mind tend to care about. After all, it’s impossible for the tokening of cognitive representations to stand in causal or correlative relations with the tokening of non-causal particulars and the existence of universals. So, SETs have to make an even bigger concession: teleosemantic explanations can only justify singling out the SEs and SE functions that ground representations of *causal particulars*.

§6. Putting the selectionist cart before the explanatory horse

This final concession reveals a deeper problem with the particular naturalization project that many friends of SETs advocate (Garson, 2022). Philosophers of mind often appeal to richer norms like truth, reasons, and rationality to account for representations of non-causal particulars and universals (Williams, 2019). These richer norms are categorically different from SE functions but they do an excellent job at explaining human behaviour. Such explanations justify singling out truth, reasons, and rationality, not singling out SEs. If SETs justify singling out SEs by appeal to explanation, then we should give priority to explanations even when they do not justify singling out SEs. Thus, we should prefer to single out these richer norms, not SEs, for explaining these representations.

³ For the record, I doubt that this is true: I think teleosemantic theories give the wrong kind of explanation for this.

However, friends of SETs often hold out for “the ambitious conclusion that multiple domains of normative facts are reducible to facts about natural selection” (Christie et al., 2022). After all, they seem to think that the only adequate, non-mysterious explanations of behaviour and other processes must be *natural* (Garson, 2022), that SE functions are *natural*, and that truth, reasons, and rationality are *non-natural*. But this puts the cart before the horse: it aims to replace (better) explanations that appeal to richer norms in domains like semantics, medicine, and psychiatry with (worse) explanations that appeal to SE functions under the typically false assumption that we already have justification to single out SEs (Kingsbury, 2022; Garson, 2022; Shea, 2022).

We should move in the opposite direction. We should take the explanatory roles that richer norms play seriously and infer that these explanatory performances justify singling out those norms—regardless of whether they are grounded in natural properties (Dewey, 2022). This approach is compatible with a modest defense of SE functions: we can take the explanatory roles that SE functions play in behavioural interactions with causal particulars seriously and infer that these roles justify singling out SE functions. But SE functions warrant being singled out (when they do) in the same way that richer norms warrant singling out: because singling them out is necessary for our best explanations—not because they are grounded in the world in a “more naturalistic” way.

Acknowledgements

I thank the APR curatorial team for their constructive criticism, which has significantly improved this commentary. I also thank Robert Lazo for his feedback on an earlier draft and Will Cailes for pointing me towards this opportunity.

References

- Christie, J., Brusse, C., Bourrat, P., Takacs, P. & Griffith, P. (2022). Are biological traits explained by their “selected effect” functions? *Australasian Philosophical Review*, 6/4: 335–59.
- Dewey, A. R. (2022). Arbitrating norms for reasoning tasks. *Synthese*, 200, 502.
- Garson, J. (2022). What are functions good for? *Australasian Philosophical Review* 6/4: 374–85.
- Kingsbury, J. (2022). Changing the subject? Christie, Brusse, et al. on the Selected Effects Account of Biological Function, *Australasian Philosophical Review*, 6/4: 367–73.
- Neander, K. (2017). Functional analysis and the species design. *Synthese*, 194/4: 1147–68.
- Okasha, S. (2022). Function in the light of frequency-dependent selection. *Australasian Philosophical Review*, 6/4: 386–99.
- Shea, N. (2022). Can a Theory of Content Rely on Selected Effect Functions? Response to Christie, Brusse, et al. *Australasian Philosophical Review*, 6/4: 400–11.
- Williams, J.R.G. (2019). *The metaphysics of representation*. Oxford University Press.