A Different Kind of Property Cluster Kind

[Draft of 10 April 2011]

Matthew H. Slater
Bucknell University
http://www.facstaff.bucknell.edu/mhs016/

1. Introduction

Richard Boyd (1988, 1991, 1999) has long campaigned for a view of natural kinds he calls the *Homeostatic Property Cluster* (HPC) account. This account has been particularly exciting for philosophers of biology unhappy with traditional essentialism about natural kinds and the views that biological kinds are, in one way or another, "historical entities". Though defenders of HPC kinds have done much to further articulate the view (see, for example, Kornblith 1993; Griffiths 1997, 1999; Wilson 1999, 2005; Wilson *et al.* 2007; Chakravartty 2007), many questions about the account remain. One pressing question concerns the way in which HPC kinds are supposed to accommodate our inductive and explanatory practices via causal, homeostatic mechanisms.

In addressing the question, I am led to an alternative conception of what a property cluster kind might be. I call my alternative, in deference to its Boydian origins, the *Stable Property Cluster* (SPC) account of natural kinds. Rather than emphasizing homeostasis or causal mechanism, the SPC account emphasizes the *stability* of a property cluster over the various ways stability may be maintained. In virtue of this modest change of emphasis, the SPC account does not advertise itself as an account of a *kind* of natural kind — it is, I think, an attractive candidate for a *general* natural kind concept, able to accommodate the diversity of natural kinds we find in the world.

I begin by motivating the cluster approach to natural kinds: how it fell initially to the neo-essentialism of Kripke and Putnam but was later reinvigorated by Boyd (§2). Because previous discussion of neo-essentialism tended to focus on the exciting semantic and metaphysical theses Kripke and Putnam proposed — some thought it a resurrection of Aristotelian essentialism — the affinity between essentialist kinds and HPC kinds is sometimes under-emphasized. There is, I suggest, a deep *pragmatic* affinity between these two accounts which the SPC account helps illuminate (§3). After outlining a few methodological and metaphysical worries for the HPC account's invocation of causal mechanisms (§4), I outline a conception of stability for property cluster kinds and explore the sense in which it underpins the reality of kind distinctions (§§5–6). The final section (§7) considers species as a case study for the SPC account, illustrating a way in

which the independently-motivated features of the SPC account do better than the HPC account at addressing worries about treating species as natural kinds.

2. The Fall and Rise of Cluster Concepts

Recall that cluster accounts of reference were seen as an improvement on what Putnam called the "traditional view" according to which the meaning of kind terms is given by a conjunction of properties (Putnam 1975, 140). On the cluster view, rather than possessing a set of necessary and jointly sufficient conditions, names were associated with a cluster of descriptions none of which were necessary — so long as a thing featured "enough" of the properties so clustered. The whole loose cluster is, in a sense, the essence of the kind. What Kripke argued so convincingly about the case of proper names (like 'Aristotle' or 'Moses') is that even this loose cluster of properties is often obviously neither necessary nor sufficient. This strongly suggested that the same difficulty would beset natural kind terms.

This intuition, however, depends on accepting that there are natural kind essences to be discovered — and there are good reasons for thinking that the discovery account of kind essence is in some trouble. In any case, one can clearly accept the direct reference theory for some types of names (those of people, say) and yet deny that it plausibly applies to natural kinds. While this is clearly a logical possibility, what positive reason can we offer for accepting the cluster view in even this narrow realm of cases?

Let's remember the motivation behind extending the direct reference account to kind terms and positing essences in the first place. As Putnam notes, natural kind terms have a special place in science: they designate "classes of things that we regard as of explanatory importance: classes whose normal distinguishing characteristics are 'held together' or even explained by deep-lying mechanisms. Gold, lemon, tiger, acid, are examples of such nouns" (1975, 139). This "holding together" helps explain several features of natural kind terms: First, that members of a natural kind have associated properties more or less in common. We would never have been tempted to a cluster approach if it had not seemed that natural kinds had more than one property in common. Second, that they might find a role in inductions and explanations: my observation that several objects possess a series of properties in common might suggest to me that I may really be observing effects of a common cause. Perhaps a subset of those effects might reliably indicate the presence of the cause and thus of a larger set of effects. Chakravartty offers a compelling metaphor:

> Properties, or property instances, are not the sorts of things that come randomly distributed across space-time. They are systematically "sociable" in various ways. They "like" each other's company. The highest degree of sociability is evidenced by essence kinds, where specific sets of properties are always found together. In

> other cases, lesser degrees of sociability are evidenced by the somewhat looser associations that make up cluster kinds. In either case, it is the fact that members of kinds share properties, to whatever degree, that underwrites the inductive generalizations and predictions to which these categories lend themselves. (Chakravartty 2007, 170)

Once we appreciate the importance of the "clustering" of properties, we naturally ask after its explanation. Essences neatly fulfill this role: the "sociability" of a kind's properties is explained by their being jointly instantiated by the possessor of the kind. Indeed, this may a good reason for counting a certain property as a kind essence in the first place.

The essentialist explanation of clustering is so good — and apparently so prevalent — that it is tempting to suppose that it is the *only* possible explanation. Devitt expressed this attitude in his argument for intrinsic biological essentialism. The law-like truth of generalizations about biological taxa (such as 'Indian rhinos have two horns') demands explanation: "There has to be something about the very nature of the group . . . that, given its environment, determines the truth of the generalization. That something is an intrinsic underlying, probably largely genetic, property that is part of the essence of the group. Indeed, what else could it be?" (2008, 352). But why suppose that there is any *one* explanation for property clustering, much less that it is the existence of essential properties attaching to the kind?

The insight of Richard Boyd's version of the cluster account is precisely that individual essences are not always needed as the glue holding together a cluster of properties — a cluster of properties "can hold itself" together. He writes:

> I argue that there are a number of scientifically important kinds (properties, relations, etc.) whose natural definitions are very much like the property-cluster definitions postulated by ordinary-language philosophers except that the unity of the properties in the defining cluster is mainly causal rather than conceptual. (Boyd 1991, 141)

Kornblith puts the central claim this way: "Natural kinds involve causally stable combinations of properties residing together in an intimate relationship" (1993, 7). The "intimacy" or "sociability" of these clusters is maintained by what Boyd calls "homeostatic mechanisms" — sometimes instantiated by a subset of the properties themselves.

Not surprisingly, there is a certain temptation to treat "HPC Kinds" as subclass of the natural kinds — perhaps one of a number of distinct kinds of natural kinds. Boyd's presentation encourages this thought, as he describes himself as identifying "a *class of natural kinds*, properties and relations whose definitions are provided not by any set of necessary and sufficient conditions, but instead by a 'homeostatically' sustained clustering of those properties or relations" (1999, 141; *my emphasis*). Wilson, Barker, and Brigandt take the cue; in their discussion of stem cells, they list a number of characteristics typically possessed by stem cells, noting that "there are excep-

tions, so that the above describes a genuine HPC kind" (2007, 218). This makes it sound as if rather than HPC kinds *tolerating* the non-instantiation of some of a cluster's properties in a member of the kind it characterizes, such lapses were in fact *required* for that kind to be one of the scruffy yet hip HPC underground. Even their title, "When Traditional Essentialism Fails", suggests that HPC is a sort of liberal fallback account of natural kinds — reserved for when things get too roudy for the tidy, conservative essentialist kinds to manage. That HPC and traditional essentialist kinds are meant to occupy disjoint territory should also be clear by reflecting on the oddity of thinking of microstructural essences as causal homeostatic mechanisms maintaining the stability of a cluster of properties.

Adopting this sort of compartmentalized stance about HPC kinds has its strategic advantages. For one, it allows essentialists and defenders of other accounts to reign more or less unchallenged in their separate fiefdoms. If a particular example of a purported natural kind fits poorly in the HPC mold, an HPCer needn't press the matter. She can shrug and admit that the purported kind may not be an HPC kind while maintaining that some kinds are. While there is nothing inappropriate about this maneuver, I think that we can do better. I am not alone. Alexander Bird also suggested that the HPC account

> can be extended to all natural kinds. The laws will explain why there are certain clusters; they will also explain the natures of those clusters — the loose and vague clusters in biology, the partially precise clusters of chemistry and the perfectly precise clusters of particle physics. Boyd introduces his idea in order to provide an alternative to the essentialist view of natural kinds. However, if I am right, the homeostatic property cluster approach can be expanded to include the essentialist view in respect of the kinds to which it applies. (2007, 210–211)

With minimal revision, HPCers can aspire to limn the contours of a more general account of natural kinds on which essentialist kinds are a limiting case.

We can get there by asking: What makes HPC kinds a particular *kind* of natural kind? Why are they not a different phenomenon altogether? The answer lies in the characteristic function of natural kinds within scientific investigation, a subject on which Boyd and Putnam apparently agree and to which we now turn.

3. Accommodating our Inductive Strategies

Boyd's account apparently allows us to accord a theory-neutral reading of the metaphor of worldly "joints". He writes:

> Kinds useful for induction or explanation must always "cut the world at its joints" in this sense: successful induction and explanation always require that we accommodate our categories to the causal structure of the world. Of course the em-

> piricist has a Humean conception of the reference to causal structure whereas the realist has a metaphysical one. (Boyd 1991, 139)

We uncover privileged taxonomic divisions, on this conception, by discovering which property-clusters are causally sustained. Boyd thus allies himself with Quine on the role of natural kinds in science: "the theory of natural kinds is about how schemes of classification contribute to the formulation and identification of projectible hypotheses" (Boyd 1999, 147). This "accommodation thesis" is at the heart of Boyd's HPC account. In this sense, the homeostatic mechanisms maintaining the sociability of a cluster of properties in HPC kinds are playing the same role in accommodating our inferential and explanatory practices as did Putnam's essences.

Moreover, both Boyd and Putnam clearly have the sense that *something* — be it an essence, a homeostatic mechanism, or some other feature of "the causal structure of the world" — must underlie or explain the sociability of the property cluster. There is convergence in this matter with some recent thought about the problem of induction. In his discussion of "explanationist" strategies for understanding inductive inference (pioneered by Harman 1965), Roger White notes an important difference between explaining our observations and explaining the truth of the generalization itself:

> The basic idea is that our support for the generalization that all Fs are G depends on how well it can explain our data, when we have observed many G Fs and no non-G Fs. But here we need to note a typically overlooked but crucial ambiguity. There are at least two explanatory tasks we might have in mind here:
>
> E1: explaining, concerning the Fs we have observed, why they are G
>
> E2: explaining why all observed Fs are G. (White 2005, 7)

Take E1: suppose we have only observed one F and found that it is G. The generalization that all Fs are Gs does not appear to explain why this one F (say, a blade of grass) is G(reen) — that is explained by the presence of chlorophyll (Peacocke 2004, 139). On the other hand, says White, E2 "is a claim about us, namely that we haven't set eyes on a non-G F. In E1 we are explaining instances of the generalization that all Fs are G; in E2 we are explaining the absence of observations of counter-instances" (2005, 8).

Peacocke calls E1 the "narrow-scope" and E2 the "wide-scope" reading of these explanations: for in E1, the explanation does not involve our observations — we merely describe certain Fs on that basis, seeking explanations for their qualities (2004, 140). Pretty clearly, if either reading is going to be relevant to induction, it will be the latter:

> Under the wide-scope reading it seems to me correct to say that when the enumerative induction is sound the evidence—namely, that all observed Fs are Gs—is really explained by the fact that all Fs are G. Here we are explaining a partially psychological condition—that the Fs that are observed are really G—by the hold-

> ing of a condition that (in at least this respect) is not psychological, that all Fs are G. (ibid.)

But to many, including Peacocke, it seems that the mere fact that all Fs are G isn't always quite enough to explain our observation to this effect — not, at least, in all circumstances. Its being a mere accident, for example, that all Fs are G robs that generalization of much explanatory potential:

> Suppose one hundred spinnings of a roulette wheel are spinnings in which the ball lands on red, and suppose we observed the first fifty spinnings. The fact that all of the hundred spinnings ended with the ball landing on red is sufficient to explain why all the fifty observations of spinnings are ones in which the ball landed on red. But an inductive inference to the fifty-first spinning that it will end with the ball landing on red is unsound. The generalization does give the explanation of our evidence, but we are not entitled to the inductive inference if we know the wheel to be unbiased. (ibid.)

What is required, argues Peacocke, is a commitment to the existence of "some condition C that explains why all the Fs are G" (141). Ruth Millikan writes along similar lines that "Clearly a concept having [rich inductive potential] does not emerge by ontological accident. If a term is to have genuine [inductive potential], it had better attach not just to an accidental pattern of correlated properties, but to properties correlated for a good reason" (2000, 17).

The advantage of the HPC account is that it offers a way of understanding the "non-accidentality" of property clustering without having to invoke microstructural essences. This flexibility allows HPC kinds to accommodate our inferential and explanatory practices in domains (such as biology) that have proved challenging to the essentialist. In the next section, I shall offer a few reasons for thinking that we can and should increase this flexibility further.

4. Concerns about Causal Mechanism

There are, it seems to me, four related (though distinct) reasons for moving away from the causal homeostatic mechanisms favored by the HPC account to a more general notion of stability for SPC account of natural kinds. The first is methodological: articulating a conception of natural kinds without making controversial metaphysical commitments about the nature of causal mechanism, homeostasis, and the causal structure of the world in general has certain advantages (cf. Strevens 2008, 7). The second is a concern about the ability of the notion of a causal homeostatic mechanism to sufficiently accommodate our epistemic practices. A third, related reason is that the conception of causal homeostasis to which the HPC account has appealed remains somewhat vaguely defined. Sometimes it appears to be used in a literal sense; at others, it seems a mere metaphor. A fourth reason is that for many ways of precisely articulating the notion of a

causal homeostatic mechanism, it seems doubtful that some purported kinds can be aptly handled by this account. This is not just the case for fundamental kinds (such as quarks and electrons), as I intimated above, but even for categories in the HPC's stronghold of biology.

I shall describe these concerns below. In each case, committed HPCers might dig in their heels in various ways or offer further specifications that evade the difficulties. That is to be expected. And it is welcome; for I am no enemy of the HPC account and am prepared to be swayed. Still, the worries give me enough pause to want to articulate an alternative. That alternative is to relinquish causal language in favor of the idiom of *stability*. Homeostatic causal mechanisms are but one way to achieve the stability a bunch of clustered properties.

4.1. Methodological Metaphysical Neutrality

If we philosophers agreed on how to understand causation, causal mechanism, and such notions, Boyd's suggestion that we should accommodate our inferential and explanatory practices to "the causal structure of the world" might be more appealing. As things stand, however, consensus about the metaphysics of causation seems rather distant. Indeed, it is not even clear in what *general* terms we should analyze causation. Shall we pursue the counterfactual approach, Mackie's INUS approach, Woodward's manipulationist approach, or what (to name just a few)? And what about the details? The corresponding debate about causal mechanism is in its comparative infancy (Machamer *et al.* 2000; Woodward 2002). This leads to a familiar, but nevertheless uncomfortable tension. On the one hand, notions of causal structure and mechanism seem sufficiently high-level for us to ignore the details of their analysis (as Boyd effectively does). On the other hand, we rightly suspect this analysis to matter. Substantial differences on the basic question of the terms of analysis likely imply substantial differences for the analysanda.

We can, of course, attempt to reduce the tension by finding a middle path — "assuming as little as possible about the metaphysics of causation", as Michael Strevens puts it about his causal-explanatory framework. Strevens notes immediately, however, that "'As little as possible' is still something, but my project is predicated on the bet that it is the nature of explanatory relevance rather than the nature of causation itself that holds the key to the broad range of questions that I ask" (2008, 5). It is an open question whether the HPCer has this sort of substantive philosophical safety net. And it is far less clear that just any old theory of causation will suit the HPC account's requirements. Singularist accounts of causation, for example, seem ill-equipped to accept *any* account of causal structure or mechanism, if by these notions one has in mind "repeating units" of causal action. Alternatively, counterfactual or manipulation-based accounts provide the basic raw material for an account of causal structure, but how should they be assembled? One obvious suggestion would be that causal structure accrues when certain kinds of causal connec-

tions are repeated and subsequently nested. But insofar as this approach already incorporates some notion of kinds, it looks unsuitable to ground a general account of natural kinds.

Of course, it's possible that the contention about causation and causal structure will wind down or that productive neutrality will be achieved. But suppose that one was a causal skeptic — either in finding it doubtful that a univocal notion of causation would be found at the root of our best sciences (Norton 2003) or that we are up to the task of ferreting out causes. It doesn't seem that such a stance should require giving up on natural kinds. Accepting some classification systems as contributing to our inferential and explanatory practices carries no obvious commitment to the existence of something that fits the description "the causal structure of the world". Though Boyd formulates his accommodation thesis in metaphysically-loaded terms, he allows that empiricist approaches might also be possible. But he reckons that these would be approaches to *causation*. I claim that putting stability at the main level an account of a natural kind offers an attractive level of neutrality. For we can say an awful lot about stability without getting into fights about how precisely to understand it.

4.2. Homeostasis and Accommodation

As I mentioned above, essences play the role of a sort of epistemic glue for traditional natural kinds. Noting that some particular has properties P, Q, and R (all standard effects of having property E — the essence of a particular kind ϕ) gives us prima facie license for inferring that the particular *is* of kind ϕ and thus possess the other properties S, T, U to which E gives rise. E offers us this license insofar as these various causal entailments are non-accidental (in some appropriately robust sense). E's being the essence of in a sense records this fact. Even when E's identity is unknown, treating some category as a kind (in this traditional sense) may be tantamount to accepting the existence of something like Peacocke's "condition C" — that there is *some* essence such which explains the stability of properties P, Q, R, S, T, U, &c. Possessing good reason for making this commitment — being justified that kind should be treated as a natural kind — is arguably often part of the background knowledge we rely upon when making inferences from we should see it as a piece of background knowledge playing a supporting role in our inference (Godfrey-Smith forthcoming).

In HPC kinds, homeostatic mechanisms — constituted by the clustered properties themselves — are supposed to take over E's epistemic role. But it is not clear that they are up to the job; not, at least, without a considerable degree of further specification. Consider an analogy. My guest room has a separate thermostat — a homeostatic mechanism *par excellence*. It is designed to maintain a consistent comfortable temperature in the room, *when it is in operation*. But when we don't have guests, we switch it off. Many homeostatic mechanisms are like this: they operate only for a time, or in some but not other contexts. But the epistemic roles of kinds apparently re-

quire more stability than the mere operation of such homeostatic mechanisms provide. One might try to achieve this stability by adding a requirement that the mechanisms maintaining cluster stability are themselves supported by further mechanisms. Change the above example a bit. Imagine that it's very important that the room remain a comfortable temperature, but my thermostat has a tendency of switching off after a few days of use. Being lazy but technologically ingenious, I design a robot to watch the thermostat and switch it back on whenever it goes out. But now we must ask after the stability of the robot-watcher's mechanism. What watches the watcher? The threatened regress can of course be stopped at any stage by offering a mechanism that guarantees the sort of stability that accommodation to our epistemic practices demands. But now we are elsewhere from a mere homeostatic mechanism.

4.3. Absence, Discord, and Conventionalism

I just claimed that homeostatic mechanisms by themselves are insufficient for accommodation. I now claim that they are also unnecessary. Many scientifically important categories are associated with clusters of properties whose stability is not plausibly maintained by causal homeostatic mechanisms. This is quite clear for fundamental particles such as electrons or quarks which apparently possess their properties (and possibly dispositions) as a matter of apparently brute nomic necessity. HPCers can handle such cases easily by simply restricting the scope of the HPC account to cover only the categories where essentialism fails. But there remain some cases apparently within HPC's purview for which it is difficult to make out the activity of causal mechanisms.

Consider sibling species. In Mayr's famous discussion of *Drosophila persimilis* and *pseudoobscura* he reports that though initially thought to be physically identical, a number of differences were eventually discovered (1963, 35). His presumption, of course (as a trenchant defender of the biological species concept), is that the discovery of the existence of homeostatic mechanisms preserving the reproductive isolation of the *Drosophila* species (and thus the stability of each species' properties), compels their basal separation. It does *not*, however, compel their separation as natural kinds: for we might very plausibly regard the whole *Drosophila* genus as a natural kind. Suppose we do: what is the homeostatic mechanism "holding together" the cluster of properties we initially identify as characteristic of that genus? Not a propensity to interbreed — for we have two reproductively-isolated species! Ereshefsky and Matthen suggest that the two separate "interbreeding structures share a common historical origin and are subject to very similar environmental pressures: this is why members of the two species are similar to each other" (2005, 6). Wilson *et al.* similarly propose that "biological individuals often are as they are and behave as they do because of the relations in which they stand" (2007, 198).

It seems a stretch, however, to regard such historical origins as *homeostatic mechanisms*. Common history and selective regimes are not obviously *mechanisms* at all — nor are they clearly homeostatic. Again, *Homeostasis* seems like the wrong metaphor here. What we have is not the resistance to disrupting a cluster of properties by the workings of certain causal connections but the stability of such clusters due to their causal *isolation* — the *absence* of potentially-perturbing causal pathways from the here and now to the there and then. Such cases seem better characterized by what Griffiths calls "phylogenetic inertia" (1999, 220) — they are "frozen accidents".

One might object at this point that causal homeostatic mechanisms underpin phylogenetic inertia insofar as reproduction and development are (relatively) high-fidelity copying mechanisms. This looks ad hoc. For it is uncontroversial that they are low-fidelity enough to allow for the variation that fuels evolution by natural selection to take place. It seems somewhat "convenient" to regard reproduction and development as maintaining the coherence of a cluster of properties just when selection and drift *do not* disrupt the homogenizing activity of reproduction and development. These processes become the relevant mechanisms unless they are not.

Once one sees this pattern, similar cases are relatively easy to spot. I have elsewhere suggested that different enantiomers of biochemical species might be considered to be distinct natural kinds (Slater 2005). But enantiomers seem like paradigm cases of kinds whose characteristic properties are maintained neither by a common microstructural essence nor causal mechanism. They are merely "stuck" in space in ways that are causally isolated and given certain types of chiral environments lead to the manifestation of different characteristic dispositions.

If the above examples are unconvincing, reflect on the fact that many of the sorts of processes Boyd and other HPCers focus on can actually be engaged in pulling a kind apart. Consider a species taxon which plays the role of a natural kind for certain questions and pursuits. Clearly, we cannot count on this category *always* playing these roles — natural selection might disrupt the stability of the associated property cluster. We need to remember that in some cases the disrupting influences of selection are ever-present. Individuals on the extreme ends of a trait parameter may have an evolutionary edge over their more moderate kin (see Figure 1). Such selective regimes are quite implausible as "homeostatic mechanisms"; nor need they be "heterostatic". But discovering that disruptive selection is operating on a particular taxon need not besmirch its epistemic utility. Accordingly, it does not show us that we were wrong to treat that category as a natural kind. For in certain scientific contexts (such as conservation ecology, medicine, functional biology) where stability across evolutionary timescales is of little concern, even species in "heterodynamic" selective regimes may possess a cluster of properties which are stable enough to afford inference and explanation in those contexts.

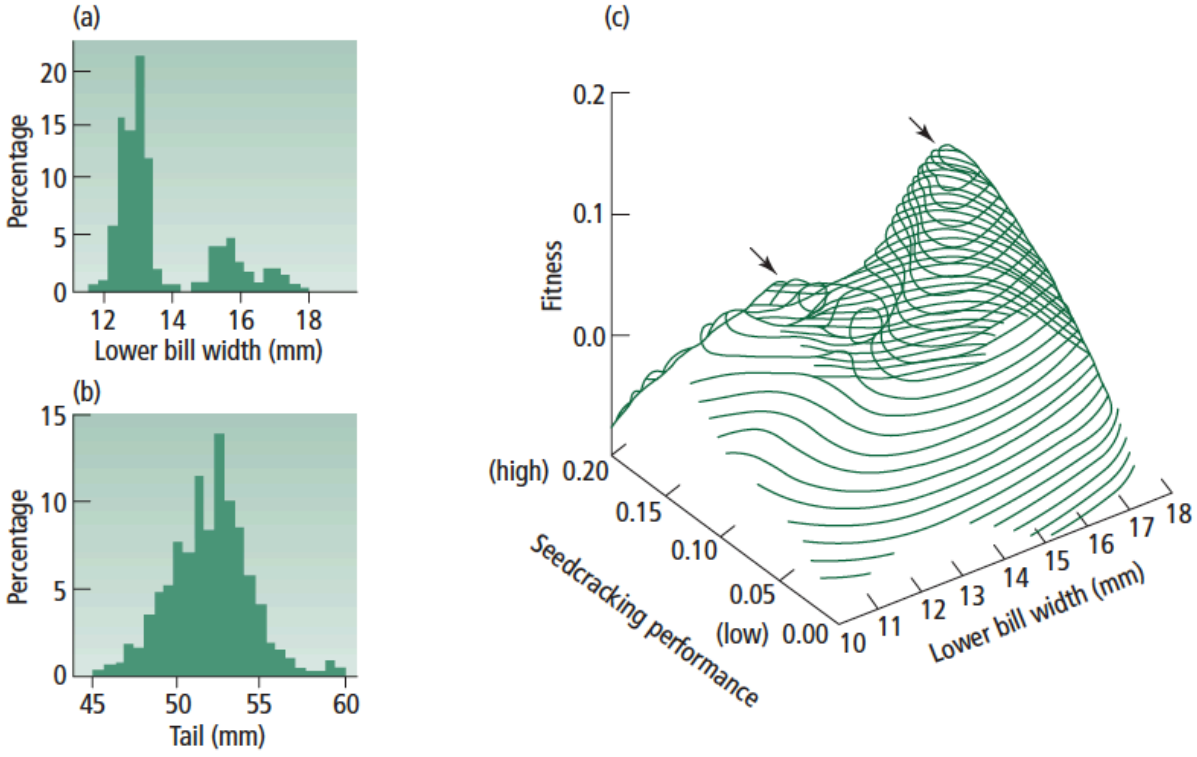

Figure 1: Disruptive selection in the seedcracking finch *Pyrenestes ostrinus.* (a) Beak size distribution; (b) General body size distribution; (c) Fitness shows "twin peaks" corresponding to the peaks and valleys in the frequency distribution in (a). (Figure and caption from Ridley 2003, 80)

It might be thought that as soon as we grant the workings of these heterodynamic forces working to pull apart clustered properties, we see right away that there must be opposing homeostatic mechanisms (intrinsic or extrinsic) at work as well maintaining the stability of the clusters. The mechanisms of reproduction and development will again seem tempting candidates. But the fact of short-timescale stability in the face of disruptive selection needn't imply the imperfect workings of homeostatic mechanisms ultimately fighting a loosing battle. Consider diamond, that highly sought-after allotrope of carbon. There is a sense in which diamonds are *not* forever: the transition from diamond to graphite in standard conditions is spontaneous ($\Delta G° = -2.9$ kJ/mol). But it occurs so *slowly* as to never be witnessed by human eyes. Only at higher temperatures does the transition occur at an observable rate (Shriver *et al.* 1994, 480). Diamond is said to be *metastable* in this sense — a phenomenon also attributed to certain organic macromolecules like proteins, DNA, and RNA. I suggest that we can identify a sense of metastability for biological taxa as well. Even in cases where populations are gradually diverging, the relevant property clusters may exhibit what we might term "biological metastability". This may in fact be the norm for species and may — depending on our interests, intentions, and such — suffice to accommodate our inductive and explanatory practices. Mimicked stability may be, for the purposes of ac-

commodation, stability simpliciter. This idea is at the core of the account of natural kinds I sketch below.

Finally, there is a more general concern about the role of causal mechanisms in the HPC account stemming from reflection on contemporary accounts of mechanism (see, e.g., Machamer *et al.* 2000; Woodward 2002). Craver (2009) has argued that the lack of objectivity about mechanisms runs afoul of HPC's accommodation thesis. In short, HPC wrongly assumes that there is an objective "mechanistic structure of the world". Specifically, Craver argues that it is not always clear whether two phenomena are expressions of the same kind of mechanism or where one mechanism begins and another ends. To the extent that one ties the identity of a kind to a particular mechanism, Craver writes, "one can be led to lump or split the same putative kind in different ways depending on which mechanism one consults in accommodating the taxonomy to the mechanistic structure of the world" (583). If Craver is correct that "human perspectives and conventions enter into judgments about how mechanisms should be typed and individuated" (591) — and I tend to think he is —, then it would appear that what natural kinds on the HPC view will depend on those perspectives and conventions. This leads to what many will regard as an unacceptably conventionalist pluralism about what kinds there are. Such a result concerns me less than the fact the notion of a causal mechanism remains vague enough as to be well-nigh useless as a starting point for dividing the world into kinds.

It might be objected that I am over-interpreting what is meant to be a metaphor. If that's so, I am better interpreted as challenging the aptness of the metaphor. Anyway, an account of natural kinds ought to rest on firmer theoretical foundation than a metaphorical similarity to other known entities and processes (especially when they appear to already presume some conception of "kindness").

Now, it may be possible to articulate a conception of causal homeostatic mechanisms that would circumvent the concerns I have articulated here. Even so, the results are likely to inherit other difficulties and controversies that have long dogged debates about causation and mechanism. Rather than contending over these particulars, it seems prudent to focus on the intended *effect* of the existence of essences and the operation of causal homeostatic mechanisms in virtue of which natural kind categories contribute to our epistemic practices: that the clustering is, in a sense to be discussed, *stable*. This is the central idea of the alternative cluster account I sketch below.

5. Stable Property Cluster Kinds

5.1. The Basic Idea

Let us return briefly to the example of the roulette wheel. Peacocke's demand that some *condition* ensure the repeated red-spinnings might be seen as a way of ensuring the stability or robustness of this occurrence. The modal fragility of such an occurrence — the sense in which had various things gone very slightly differently with the spinning, we wouldn't have seen fifty reds — seems to imply a *temporal* fragility that cancels our inductive warrant. This seems quite plausible. If some truth could easily have been false, then who's to say that it won't *go* false any time from now. On reflection, however, the general implication does not hold. There are lots of facts that are modally fragile in the sense that they very easily might not have occurred — for example, that Cletus the clumsy archer won the archery tournament yesterday — but which are not in any danger of *ceasing* to be true in the future.

Even if one rejects this modal-to-temporal implication, one might still insist that our inductive and explanatory practices need to be grounded by more than just dumb luck. However this may be, I see little reason to focus on the various precise ways in which the stability of a cluster of properties obtains. What matters for the epistemic utility of kind categories is not the cause or explanation of their stability, but *the fact that they are stable*. Peter Lipton expressed a similar sentiment when, in commenting on Kornblith's application of the HPC account to the problem of inductive knowledge, he wrote: "Essences are supposed to hold together observable properties in stable clusters, but it is not made clear why this should make for a more inductively knowable world than one where that stability is a brute fact" (1996, 493). Accordingly, the Stable Property Cluster (SPC) account of natural kinds identifies natural kinds (for particular sciences) as categories associated with clusters of properties (in the HPC sense) which are sufficiently stably coinstantiated to accommodate the inferential and explanatory to which those sciences put such categories.

This subtle shift of focus scores three significant goals. First, the SPC account evades the problems mentioned above. Even categories associated with clusters of properties whose sociability are gradually being disrupted by heterodynamic selection regimes can underpin our epistemic practices in virtue of their having sufficient stability for the purposes of the relevant sciences. Ditto for categories whose cluster's stability is not maintained by any mechanisms in particular. Second, it achieves an attractive degree of neutrality. Stability, as I shall understand it is a high-level concept that is independent of its particular realizers and their analysis. Third, it represents a more general account of natural kinds able to encompass strict essentialist kinds, kinds with historical essences (Griffiths 1999; Okasha 2002), and HPC kinds whose clusters' stability are maintained by genuinely homeostatic causal mechanisms. The stability that lends itself to a kind's inductive and explanatory utility is, as philosophers of science are apt to say, multiply realized.

5.2. Two Conceptions of Stability

What might it mean to call a cluster of properties "stable"? One first stab might focus on the instantiation of the clustered properties by a particular individual — an instance of the kind associated with that cluster. Say that a property cluster kind ϕ is *instance-stable* when (to a first approximation) satisfiers of ϕ (individuals of the kind ϕ) do not easily relinquish the relevant cluster of properties. When instantiated, the instances of the properties in resist their non-instantiation — perhaps by constituting a homeostatic mechanism.

Instance stability is both too strong and too weak for characterizing natural kinds. It is too strong inasmuch as it implies that kind membership is "sticky". Once a particular thing satisfies the cluster of properties associated with a kind ϕ, it resists becoming not-ϕ. But while this may be characteristic of some kinds (cases where kind-essence and individual-essence are somehow bound together), other objects apparently change their kind quite readily. Instance stability is too weak in that it does not sufficiently account for kinds' epistemic role. To see this, let us consider a schematic example.

Suppose that ϕ is a kind associated with a cluster of properties *P*, *Q*, *R*, *S*, *T*. We can use brackets to denote that these properties are clustered without presuming the involvement of other abstracta (to wit, sets): [*P*, *Q*, *R*, *S*, *T*]; for convenience, let us also use the name of the kind enclosed in brackets as a notational equivalent: [ϕ]. We can think of these as predicate in the normal way — e.g., *j* is a member of cluster kind ϕ could be symbolized as '[ϕ]*j*' or '[*P*, *Q*, *R*, *S*, *T*]*j* ' (depending on how specific we wished to be) — so long as we do not think of this as reducing to either '*Pj* ∧ *Qj* ∧ *Rj* ∧ *Sj* ∧ *Tj*' or '*Pj* ∨ *Qj* ∨ *Rj* ∨ *Sj* ∨ *Tj*'. For on the cluster kind view, for *j* to be a member of is not for it to have all of (or at least one of) *P*, *Q*, *R*, *S*, *T*. Rather, it is for *j* to have a *cluster* of those properties. I shall have more to say later about what constitutes a cluster; for the meanwhile, I leave it at the intuitive level — something along the lines of "*a goodly many of the properties in question*". We can represent 's instance-stability as the claim that for all *x*,

$$[P, Q, R, S, T]x \rightarrow \blacksquare[P, Q, R, S, T]x$$

using '■' as a sort of "robustness" operator, temporal or modal: the precise interpretation of this operator is not my present focus. But ϕ's being an SPC kind in this sense does not do justice to the inference pattern mentioned in §4 above, where the observation that *j* has *P, Q, R* gives us good reason for expecting that *j* has S and T as well. What we want out of clusters is not mere "sociability" — that once a cluster of properties are together instantiated, they are hard to scatter —, but *cliquishness*. Peg, Quinn, Ralph, Sarah, and Tim form a clique, say. Spotting Penny, Quinn, and Ralph at the mall means that Sarah and Tim are probably there as well. Nothing is

implied about how long they'll stay. Perhaps they flit from place to place, but when a few of them are around, you can bet that the others will be as well.

Call this conception of stability *cliquish stability*. This is a rather more "abstract" variety of stability: a cluster of properties can be cliquishly-stable without its being instance-stable. The idea is to capture the fact that some properties are clustered such that possession of some reliably (if imperfectly) indicates the possession of whole cluster (if not each property in the cluster). To be more precise about this, we need the notion of a "sub-cluster". Consider again our property cluster [ϕ] (that is, [*P*, *Q*, *R*, *S*, *T*]). Let a sub-cluster of [ϕ] include some but not all of the properties in [ϕ]. 'Some' here is meant to be interpreted not as the familiar existential quantifier of first-order logic, but as the *more* familiar colloquial usage as an indefinite plural quantifier (as in "some philosophers got inebriated at the Smoker"). [P, R, T] would thus be a sub-cluster of [ϕ]; as would [Q, R], and so on. For convenience, assign arbitrary names to these sub-clusters by simply subscripting [ϕ]: [ϕ_1], [ϕ_2], [ϕ_3], ... Now let us say that a property cluster [ϕ] is cliquishly-stable when for all x and for many sub-clusters [ϕ_1], [ϕ_2], [ϕ_3],…:

$$\blacksquare(([\phi_1]x \Rightarrow [\phi]x) \wedge ([\phi_2]x \Rightarrow [\phi]x) \wedge ([\phi_3]x \Rightarrow [\phi]x) \wedge \ldots)$$

where the '$\Rightarrow$' is to be read as probabilistic entailment. And again we have the "black box" of robustness. What to say about this?

The option I'd like to explore here takes its cue from Marc Lange's treatment of laws as members of a certain type of stable set. The basic idea is that certain sets of truths are maximally invariant under counterfactual perturbations. Given any counterfactual consistent with the members of the set, had it been the case, then all the members of the set would still have been the case. Ditto for an arbitrarily nested sequence of counterfactuals: had anything at all compatible with the set been the case, then had anything else compatible with the set been the case, (and so on), then the members of the set would still have been true. Lange calls this special kind of stability "non-nomic stability" (for extensive discussion, see Lange 2000, 2009). Take some logically-closed set of truths Γ. This set possesses non-nomic stability if and only if for each member m of Γ:

$$p \mathrel{\square\!\!\rightarrow} m,$$

$$q \mathrel{\square\!\!\rightarrow} (p \mathrel{\square\!\!\rightarrow} m),$$

$$r \mathrel{\square\!\!\rightarrow} (q \mathrel{\square\!\!\rightarrow} (p \mathrel{\square\!\!\rightarrow} m)), \ldots$$

for any non-nomic claims p, q, r, … which are logically consistent with the members of Γ. Lange conjectures that the set of laws is the only non-trivially non-nomically stable set. This affords a sharp distinction between the facts which are laws and those which are accidents.

However, Lange allows that there are different sets which are stable on more restrictive ranges of counterfactual suppositions. That this is so is especially clear for certain biological sciences. Some biological generalizations, though they clearly could have been false, possess a significant degree of stability in the face of these more restrictive ranges counterfactual suppositions. Consider an example. Lange mentions the belief among anthropologists that "any person of entirely Native-American heritage is blood type O or blood type A" (2000, 13). Though a historical accident — "research has suggested that all Native Americans are descended from a very small band that crossed the Siberia-Alaska land bridge, and as it happened, allele B was not represented in that company" (*ibid*.) —, that accident features a broad range of counterfactual stability. It would still be the case had a very wide array of facts been different. To modify another of Lange's examples (2004, 106), doctors might report that a certain Native-American patient would still have gone into anaphylactic shock if the transfusion of type B blood had been administered sooner, or administered along with a different concentration of saline, or what have you. Of course, the blood type generalization probably wouldn't still be true had, say, the winter of 10,273BCE been slightly warmer — at least, we shouldn't evince much confidence that it *would*. The point is that if you are an emergency room doctor, the blood type generalization is stable enough for you to rely on; for it is unlikely that you would be very interested in counterfactuals involving the weather twelve thousand years ago. What matters to you is what manipulations in the here and now might save your patient.

Lange suggests we think of such restrictions of the range of counterfactuals under which certain generalizations are stable as being defined by the "interests" of the relevant special sciences. Change the above example slightly; Lange claims that

> it is of medical interest to know whether a given heart attack might have been less serious had epinephrine been administered sooner, or had the patient long been engaged in a vigorous exercise regimen, or had she been wearing a red shirt, or had the Moon been waxing. But it is not of medical interest to know whether the heart attack might have been less serious had human beings evolved under some different selection pressure. A physician might blame a patient's untimely death on her smoking, but not on human evolutionary history. (Lange 2004, 107)

While I am very much sympathetic to this basic idea — there are certainly certain counterfactual antecedents which are of perennial interest to certain fields and many more others which are *not* —, there are some pressing concerns about the details of how interests apportion modal space. Set these for the time being and focus on the applicability of the basic idea to SPC view of kinds.

Suppose we understand the black box in the above definition of cliquish stability in a broadly Langian way: that a property cluster [ϕ] is cliquishly-stable when for all x and for many subclusters [ϕ_1], [ϕ_2], [ϕ_3], ...:

$p \mathrel{\square\!\!\rightarrow} ((\lbrack\phi_1\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge (\lbrack\phi_2\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge (\lbrack\phi_3\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge \ldots)$,

$q \mathrel{\square\!\!\rightarrow} (p \mathrel{\square\!\!\rightarrow} ((\lbrack\phi_1\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge (\lbrack\phi_2\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge (\lbrack\phi_3\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge \ldots)$,

$r \mathrel{\square\!\!\rightarrow} (q \mathrel{\square\!\!\rightarrow} (p \mathrel{\square\!\!\rightarrow} ((\lbrack\phi_1\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge (\lbrack\phi_2\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge (\lbrack\phi_3\rbrack x \Rightarrow \lbrack\phi\rbrack x) \wedge \ldots), \ldots$

where p, q, r, . . . meet the following conditions:

(a) They are consistent with the probabilistic entailment relationships from sub-clusters to clusters;

(b) They are consistent with the natural laws (i.e., no counterlegals);

(c) They meet the relevant applicability standards.

The justification of conditions (a) and (b) is fairly straightforward: (a) is isomorphic to Lange's requirements for p, q, r, . . . in his definition of non-nomic stability. We cannot expect some fact to be stable on the assumption of its negation: this would be tantamount in this case to insisting that something remain a natural kind even if it failed to be a natural kind! The justification of (b) is likewise parasitic on Lange's construction and assumes that the laws are at least partly responsible for facts about the relationships among a cluster's properties. Actually, (b) is stronger than we really need, since there will be some counterlegals on which certain cluster relationships continue to hold. But this overkill affords some simplicity and does no harm that I can see.

Condition (c) is the interesting — because open-ended — condition. Here we apparently return to the worries about Lange's invocation of the interests of certain special sciences. One worry is that if interests play a significant role in the definition of stability, then we immediately foreclose on the "naturalness" of a system of natural kinds defined in these terms.

Another pressing question I already gestured toward: What exactly defines a discipline's interests? Is it credible that disciplines implicitly single out certain ranges of counterfactual antecedents for consideration of whether certain other facts would remain true on their supposition? (I take it for granted that they do not *explicitly* do so.) *Prima facie*, it seems far easier to make sense of certain counterfactual antecedents *not* being of interest to a particular discipline. Were you to ask one of the doctors whether she was interested in whether the heart attack would still have been as severe if the humans had followed a different evolutionary trajectory, the answer surely would be "no". But asking the doctor whether she is interested in what the scenario would have been like under different fashion or lunar conditions doesn't seem likely to elicit a different response. One might claim that the reason our doctor won't capitulate to being interested in the phase of the moon is that she already knows (or judges with a high confidence) that the moon's phase has (practically) nothing to do with her patient's heart condition. Perhaps this is enough to qualify the moon's phase as of medical interest. This strikes me as a rather odd thing to say — but perhaps only for Gricean reasons.

I prefer to think of condition (c) above as being defined in terms of "relevance" rather than interests. Interests do play a role in determining which counterfactuals are relevant to the evaluation of cliquish stability, but only in a derivative way. To illustrate, consider an example. Medicine is interested in saving lives, let's say. The inferential and explanatory work of the medical sciences focuses on this goal by considering circumstances that are within our power to control: hence the relevance to medical practice of the counterfactual suppositions involving different amounts of drugs, exercise, or time in the ambulance; and hence the *irrelevance* of suppositions involving evolutionary contingencies or circumstances likely to bear on them. It is not simply that such "frozen accidents" are without our power to change — for nor is it within our power to change how long a certain patient *spent* in an ambulance once they are at the hospital. Rather, the accidents of evolutionary history are not similar in informative ways to circumstances which *are* manipulable. In contrast, reflecting on whether a certain patient would have lived had he arrived at the hospital sooner, *is* potentially instructive for future cases. This is why shirt color and lunar phase are also relevant (in the present sense) to medical concerns. They are circumstances which are similar to those we can either manipulate (shirt color) or at least be sensitive to (lunar phase).

Property clusters which are cliquishly-stable for a given science, project, research program, or what have you thus offer certain fixed points for those inquiries in the sense that for possible manipulations relevant to those pursuits we may count on finding the clustered properties together, where we find them. This raises the possibility that some clusters are only natural kinds for particular domains of inquiry. The result is straightforwardly parallel to Lange's treatment of special science laws, but not uncontroversial for this reason. In particular, one might worry that this would squash any hope of the SPC account offering a sense in which kinds "carve nature at its joints".

6. Interests and Realism about SPC Kinds

SPC kinds can be, as I will put it, "*domain-relative*" in several senses. First, there is the question of how many properties are expected to be clustered together for something to count as a kind. Second, the norms and aims of certain domains may require different levels of cluster cohesiveness — that is, different disciplines may tolerate different degrees of flexibility in the clustering required by their respective kinds. Perhaps property clusters defining physical kinds like electrons or quarks are supposed by those disciplines to be perfectly clustered (conjunctive) kinds while those of higher biological taxa like families may be quite loose. A third, closely related, sense in which cluster kinds can be domain-relative is in how the probabilistic entailment relations are understood: how likely is it that the instantiation of a certain sub-cluster betokens the instantiation of the whole cluster? Fourth, and finally, we have the interest-informed relevance condition (c) above circumscribing how we think of a cluster's stability. Inasmuch as a particular

cluster of properties can meet the requisite conditions for one domain but fail to meet them in another, we must allow that at least some collections of things only instantiate natural kinds from the perspective of particular sciences, or to pick a more neutral term: *domains*. Call this *domain-relativity*. If some kinds are domain-relative, the question of what kinds there are *tout court* is not tractable. What we can legitimately ask instead is what kinds various domains of inquiry in fact recognize or *would* recognize given their present aims and interests.

A different relativity about natural kinds may attach to certain physical contexts: call this *context-relativity*. The point can be easily made using the language of homeostatic mechanisms, but applies *mutatis mutandis* to SPC kinds. If certain mechanisms only successfully maintain the stability of a cluster in particular contexts, then such clusters fail to be natural kinds unless relativized to those contexts. Of course, contexts where the properties in the cluster typically go uninstantiated *anyway* will not generate this effect. Remember: we are not talking about instance stability. Many instances of a kind lose their properties in some circumstances (e.g., a protein denatures outside of its usual temperature range). The context-sensitivity at issue here is a more subtle affair: where the property cluster may continue to be instantiated, though the mechanism(s) maintaining its stability no longer function. For example, the cluster of properties defining certain cell types may only be stable when considered *in vivo* (rather than in a Petri dish, say).

In allowing the interests and norms of a discipline — or even a particular research project — to influence whether a certain category counts as a natural kind might seem like a rather hefty dose of pragmatism to swallow. No doubt some will be happy about this. Others may see it as a bridge too far: a theory of natural kinds, they will insist, should tell us about the objective divisions in the world that pre-exist our classificatory activities. Otherwise, we cannot make sense of some theories doing better than others at "carving nature" nor how certain schemes of classification can be in error.

This concern is overstated. Though there's a clear sense in which we cannot be mistaken about what norms and interests to adopt concerning certain classifications (since norms are not truth-apt), I think we *can* come to see ourselves as having been *wrong* to hold those norms. Perhaps we held them in the first place due to some genuine mistake (taking some homogenizing effect as more important than it actually is, say). And we certainly can be mistaken about whether a certain category is associated with properties that stably cluster *given* the relevant norms and interests (assuming that there are facts about property instantiations and what subjunctive claims which are true independently of us). Moreover, it seems plausible that there might be other ways of evaluating our norms. We may, for example, find that *failing* to relativize the evaluation of a certain cluster's stability forecloses on certain inferential advantages flowing from recognizing a category as a context-relative kind. This, of course, presumes other "meta-

norms" — such as: *whenever possible, we ought to revise our norms governing property cluster stability evaluation so as to increase our recognition of kinds*.

So while a domain's norms and interests are relevant to what natural kinds there are, it's not the case that we can arbitrarily "define nature's joints into existence". Nor do natural kinds await our classificatory activity in order to come into being. For the norms and interests relevant to a cluster's stability often pre-exist those activities. They do not, however, pre-exist us. Thus a critic might point out that there is no live sense in which the SPC account is a realist account of natural kinds — for it seems that there is no sense in which there were SPC kinds before science came on the scene a few thousand (or a few hundred) years ago. But surely there were different natural kinds of things!

I think that there are two compelling things to say in response. First, it is not clear to me that we should be aiming for a realist conception of natural kinds to begin with. I put a higher priority on maintaining some of the key realist intuitions about classification — that we can be mistaken about our systems of classification, that we can classify things in better and worse ways, and so on. Second, one may hold out for a sense in which *certain* special categories — electrons, say — are natural kinds in a norm-neutral way. I think I can assent to the spirit of such a request, if not the letter, with minimal retraction of what I've been pressing by (so to say) crossing Mill with Whewell. Recall Whewell's (1840) much cited idea about the naturalness of a system of classification stemming from the convergence of different systems of classification on the same categories. Though neat, this idea won't help us make sense of electrons having discipline-independent objectivity, since the non-physical sciences (for the most part) do not have a great deal of truck with fundamental particles. Here, however, we might invoke another nifty idea from John Stuart Mill (1872) about objects as "permanent possibilities of sensation". Perhaps there are some clusters of properties such that no matter *how* a discipline adjusted its norms and aims (compatible with the discipline maintaining an understanding of the natural world in view), the category that cluster described would be fit to play a robust epistemic role in the discipline. We might say that such categories exhibit a "permanent possibility of Whewellian convergence". This allows us to see that the pluralism resulting from the SPC account's domain- and context-relativity need not extend to *all* kinds. **[I'm very unconfident about this idea: perhaps I should drop it as orthogonal to the rest?]**

While this goes some ways toward accommodating realist intuitions, I readily admit that the SPC account exhibits some distinctively non-realist features. The context- and discipline-relativity of some kinds show that natural kinds (on the SPC view) are not an ontological category (cf. Lowe 2006). Nor is it obvious how they could be reducible to facts in other ontological categories — universals, for instance. Bird suggests that:

> a simple reduction of kinds to combinations of universals is available along the lines proposed by Armstrong. . . . [This] strategy can be extended, by considering kinds as homeostatic property clusters. Although Boyd does not see that latter in ontological terms, we can construe them as sums of properties, just as complex particulars are the sums of their component parts. (Bird forthcoming, 11)

Yet treating a cluster of properties as a *conjunction* trades away what is arguably most distinctive about the approach: its looseness and corresponding ability to accommodate the messy patterns of biological variation and similarity. The conjunctive approach is far too strong.

Richards makes essentially the opposite mistake by understanding property cluster accounts of kinds as *disjunctions* of properties: "essences [on the HPC view] are a disjunction" (2010, 154). This approach is inappropriately weak. Suppose we take a kind K as being defined as being P ∨ Q ∨ R. Two individuals could possess this disjunctive property while sharing neither P nor Q nor R in common. Disjunctive similarity is dirt cheap. The cluster approach requires more. What, then, *are* SPC kinds if not conjunctions of universals? Rather than recognize a *sui generis* category — of *clusters*, say —, I prefer to think of being a natural kind as a kind of *status* things or pluralities of things in various ontological categories can have. As we shall see in the next and final section, this approach offers what seems to me some welcome respite from a longstanding debate about the metaphysics of species.

7. Case Study: Species

[This section is very much "in progress" and will be developed further — with specific examples and philosophical issues indicated below.]

It's time to see the SPC account at work. Since species have already featured prominently as a test case for the HPC account of species [**references**], I will concentrate my discussion on some of the novelties. We've already seen reason for thinking that concern over the details of causal mechanisms are likely to affect the application of the HPC view to species. Obviously the SPC account avoids those difficulties. But *prima facie*, it inherits other difficulties detractors have raised for the HPC account. In this section, I shall focus on on in particular stemming from the fact that many species are highly polytypic that has been developed most explicitly by Ereshefsky and Matthen (2005).

While HPCers have things to say in response, it seems to me that they underestimate the force of the problem (indeed, it strikes me that Ereshefsky and Matthen do not fully appreciate its import). I propose a "bent solution" that flows from the anti-metaphysical stance taken above. Because SPC kinds are not properly speaking an ontological category, there is no pressure to attempt a univocal application of the SPC account to species. While *many* species taxa are argua-

bly SPC kinds, I believe that we ought to admit that some are not (or are "less so"). This coheres, I think, with our epistemic treatment of these taxa.

7.1. The Challenge of Species Polymorphism

"The Buckets"

Ereshefsky and Matthen claim that HPCers like Boyd (1991, 1999) and Millikan (1999) "misidentify the phenomenon" which the HPC theory attempts to solve.

> They think that natural kinds, including biological taxa, are united by similarity within the species-population, and this is what needs to be explained by homeostatic mechanisms. And it is undeniably true that there is similarity within biological taxa. Nevertheless, it is equally true that many taxa are characterized by stable and persistent differences. (2005, 7)

Familiar examples — of sexual dimorphism or differences in developmental stage — show that if species are natural kinds, natural kinds are not united exclusively by similarity relations. Hence as soon as one starts to seriously consider the prospect of characterizing species — even roughly, imperfectly — by reference to a cluster of shared properties, one encounters a biological accounting nightmare. Far from homeostatic mechanisms maintaining the similarity of members of a species, there may be mechanisms preserving a population's polymorphism: what Ereshefsky and Matthen refer to as "heterostatic mechanisms" (14).

Two knee-jerk responses to this problem should be set aside (one more gently than the other): on the one hand, as we've noted, polytypism drives many to "deep" structural properties (like genetic essences) that even members of different sexes or morphs necessarily share. This wrongly assumes that there are (and will always be) suitably deep properties to do the job. Failing this, biologists may go historical: phenetic character (either at the macro or micro level) doesn't matter nearly as much as descent. While this latter option has more going for it, it seems not the only solution. Do not HPCers have a way out?

Boyd recognizes the challenge posed by polytypic species, but sees it as requiring only elaboration of the HPC account. We may have to

> characterize the homeostatic property cluster associated with a biological species as containing lots of conditionally specified disposition properties for which canonical descriptions might be something like, "if male and in the first mold, P," or "if female and in the aquatic stage, Q" (1999, 165).

Matthen and Ereshefsky worry that this sort of maneuver will "become a universal solvent that makes all variation disappear and collapses the entire biological domain into a single morpho-clump", thus trivializing natural kinds. "One could then regard the vast network of ecological relationships that constitute the biosphere as a single homeostatic property cluster maintaining polymorphism in 'Gaia'" (2005, 9). This seems an overreaction (Wilson *et al.* 2007, 210–211). That we *can* identify terrestrial life as of a kind does not imply that we *ought* to (ignoring nested kinds in the process). Whether HPCers are right to be nonplussed by the status of Ereshefsky and Matthen's concern, there are some deeper problems with the conditional response to the polymorphism problem.

First, while many kinds of polymorphism occur in connection to independently-specifiable physiological conditions (e.g., sex or developmental stage), it is by no means clear that all polymorphism can be. In some species, adults of each gender exhibit striking polymorphism [**refs**]; in others, different environmental contexts trigger the disparity. Just take the most familiar polytypic species: dogs. What conditional size specification might we attach to Great Danes and Chihuahuas? Employing those names (as in 'If Great Dane, then …') looks like it would incur a heavy dose of circularity.

More importantly, a second problem concerns how we should interpret these conditional properties. It's difficult to tell what Boyd's intentions are from what he writes. While he mentions dispositional properties, he speaks of conditionals — many doubt that these are equivalent. Set that worry aside for the moment. Those who think that dispositional properties are *reducible* to conditionals, typically finger subjunctive conditionals for this task; yet Boyd's examples appear to be indicative conditionals. There are difficulties with either interpretation.

Suppose we elect the latter interpretation. Consider a certain species S which features a tidy sexual dimorphism. Members of S (typically) share the complex logical property 'being an x, such that (if Mx, then Px) and (if Fx, then Qx)', where M and F stand for the sexes and P and Q, other biological properties not shared between males and females. This logical property is conjunctive in character. Accordingly, it can be decompose into two components which are conditional in character. So far so good. But conditionals are equivalent to disjunctions. So our initially unproblematic-looking logical property boils down to a pair of *disjunctive* properties: 'either not being male or being P' and 'either not being female or being Q'. And as we noted earlier,

there are compelling reasons for thinking that disjunctive properties do not count at all toward similarity.

The subjunctive reading of the relevant property avoids this difficulty, but also raises awkward questions. Take a particular male of S; is it the case that if it was female, it would be Q? That's not clear. As P.D. Magnus (MS, 5) points out, performing a sex-change operation on a duck won't generally change its plumage. Moreover, there's something suspicious about attributing to something a subjunctive property grounded in the nature of *other* things. If I were a Tea-Party Republican, I would be against universal healthcare. What makes that statement true is not a fact about *me*, but a fact about Tea-Party Republicans. One might fairly doubt the aptness of this analogy. Perhaps there's some sense in which the conditional properties of polytypic species *are* imminent in them — say, in their genetic code. But notice how perilously close this move gets us to a genetic form of essentialism. The property in question is no longer a conditional property spelled-out in general terms, but something resembling a genetic essence. This appears to amount to a dilemma. Insisting that the subjunctive conditional properties are grounded in the intrinsic nature of individual members of S runs afoul of anti-essentialism; *not* insisting on this incurs the suspicion that just about *anything* could possess that property. If my dog Mabelle was a Tea-Party Republican, she too would reject universal healthcare! Ereshefsky and Matthen's "universal solvent" problem returns with a vengeance.

Return to the suggestion that members of polytypic species share *dispositional* properties. Insofar as this approach inherits the difficulties of mentioned immediately above, it doesn't appear to be much help. Nor does treating dispositions as irreducible "powers" (Molnar 2003) look any more promising. Dispositions are typically understood as properties associated with a particular kind of manifestation event (e.g., shattering) triggered by certain sorts of stimuli (e.g., striking sharply). Whether or not these properties are analyzable or reducible to these terms, it doesn't seem that "polymorphic properties" have this character. Certain species of ducks are disposed to feed at the beginning and end of the day and to take flight when alarmed. It seems odd to say that they are disposed to have green heads when male. And in any case, the worry above about the apparently extrinsic ground of such facts appears to infect treating them as dispositions (see Magnus MS, 4–5 for discussion).

7.2. Are Species Real?

We have seen so far that straightforward responses to the polymorphism problem for cluster accounts of species face some difficulties. My strategy will be to evade polymorphism problem, rather than face it head on. One of it crucial assumptions — which I am happy to abandon — is that highly polytypic taxa must be treated as natural kinds by the HPC (or some other) account of natural kinds. Denying this effectively exempts kind-advocates from having to account for the

polymorphism within those taxa. This move parallels the dialectic concerning the ontology of species thirty years ago: when advocates of essentialistic accounts of natural kinds finally faced up to the heterogeneity of species at the genetic level, it did not trouble *essentialism* about natural kinds — just the thesis that *species* were natural kinds. This, of course, was one of the key stimuli of the rise of the species-as-individuals thesis.

There is an obvious objection to this line of thought: *There is an important disanalogy between the historical case and the present one. In the historical case, the denial that species were natural kinds was part and parcel to a univocal approach to their ontology. Denying just that polytypic taxa are natural kinds looks comparatively ad hoc. Further, it seems to relinquish one of the key motivations for treating species as natural kinds: to make sense of the intuitions that they are (in some sense) real features of the world. Some species taxa are polytypic. If we allow that such taxa are not natural kinds, avoiding the polymorphism problem comes at the price of denying that species are natural kinds.*

There is an ambiguity in this analysis, however. What is it for someone to claim that species are natural kinds — or real, for that matter? One interpretation involves a claim about the ontological category of all species taxa as part of a top-down claim about the species category — that any taxa at the species rank must have such and such a metaphysical character. On this interpretation, it makes perfect sense to ask (after Ruse 1987) "Are species natural kinds, individuals, or what?" and expect a univocal answer. A second, weaker interpretation treats the question as focussed on species taxa without the presumption of univocality. Asking whether species are natural kinds, on this interpretation, is tantamount to asking whether *some* species are natural kinds.

The most compelling reason I can see for preferring the first interpretation over the second stems from treating the issue as metaphysically-meaty question about what ontological category species taxa should fit into. Since I deny that SPC kinds *are* an ontological category, I tend to prefer the weaker interpretation of the question on which some but not all species taxa might count as natural kinds. This might seem disappointing at first (reality often is!), but I see it as simply reflecting the fact that not everything we identify as a species taxon in fact play the epistemic roles I believe are characteristic of natural kinds. But it should hardly be surprising, given the diverse justifications for treating certain populations of organisms as natural kinds (not all of which focus on their epistemic utility). It is for this sort of reason that I am deeply skeptical of claims like Devitt's "that Linnaean taxa have essences that are, at least partly, intrinsic underlying properties" (2008, 346). Whatever one thinks of essentialism about natural kinds in general, if this is taken to mean that *all* Linnaean taxa have such essences, it makes us too good at discerning such taxa. Compare: suppose that, incredibly, biologists discover that in fact "Bigfoot" is real — that is, there is a very reclusive upright primate species striding through isolated woods in North America. Clearly this is compatible with many of Bigfoot "sightings" being the product of

confusion, mistake, wishful thinking, or deliberate hoax. Likewise, I believe that we can be right to say that species are real in virtue of there being *many* species taxa which are SPC natural kinds.

I cannot offer a full explication of the connection between being an SPC kind and being "real" in this context — other than to concur with much of what Boyd has to say (mutatis mutandis) about HPC kinds being real to the extent that they are fit to serve a key role in our epistemic practices (1991, 139). Hopefully it seems clear enough that playing such a role is at the bottom of many of the strongest intuitions for thinking that (many) species taxa are real in the first place. Whether such intuitions can be fully vindicated is a difficult question. I think we should be prepared to settle for somewhat less.

So the following sort of "dappled" picture of species taxa emerges. Some collections of organisms share a robust cluster of properties in common which, in virtue of that cluster possessing cliquish stability, are apt to serve in characteristic inferential and explanatory roles in particular disciplines and contexts. In some cases (and perhaps for some contexts and purposes), the polymorphism is swamped by the cluster similarity which ignores the strongly varying parameters. In other cases, the polymorphism is deep enough to trouble our inferences at the species level. We simply would not infer to ethological conclusions (say, about breeding behavior) from the fact that a particular organism was a mallard. We *would*, probably, make such inferences from the knowledge that that organism was a *male* mallard. It might happen that populations of particular sexes, developmental stages, subspecies, ecotypes, and so on within a species count as SPC natural kinds but that the whole species does not (or does only to an inferentially weak extent). Such species might be viewed as genealogically-defined assemblages of natural kinds. I take this suggestion as analogous to Sterelny and Maclauren's view of species as "ecological mosaics": "ensembles of populations, each with its own niche" (2008, 38; see also Sterelny 1999). **[Need to explore an example here in more detail, I think.]**

What the SPC account of natural kinds offers us is a flexible, high-level approach to understanding the various ways in which different categories can be regarded as genuine features of the world. Clearly many questions about the approach remain. I will close by mentioning two that stand out as especially urgent. First, how in detail should the intuitions that SPC kinds are "genuine"/"real" features of the world be squared with the various sorts of relativity I mentioned above? Second, how should we understand the metaphysics (and epistemology) of cliquish stability (even within a particular context)? By taking a page out of Lange's book, should we also follow him to primitivism about subjunctives or propose some different account of what makes subjunctives true? I would prefer to not take a stance about this difficult question, elevating the concept of cliquish stability to a high-enough theoretical level to avoid the fray below (in some-

thing like the manner of Lange 2005), but doing so may be unavoidable. In any case, these are matters for another occasion.

REFERENCES

Bird, Alexander (2007) *Nature's Metaphysics: Laws and Properties*. Oxford: Oxford University Press.

——— (forthcoming) "Are Any Kinds Ontologically Fundamental?", in Tuomas Tahko (ed.), *Contemporary Aristotelian Metaphysics*, Cambridge: Cambridge University Press.

Boyd, Richard (1988) "How to Be a Moral Realist", in Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism*, Ithaca: Cornell University Press.

——— (1991) "Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds", *Philosophical Studies* 61:127–148.

——— (1999) "Homeostasis, Species, and Higher Taxa", in Robert A. Wilson (ed.), *Species: New Interdisciplinary Essays*, Cambridge: MIT Press.

Chakravartty, Anjan (2007) *A Metaphysics for Scientific Realism*. Cambridge: Cambridge University Press.

Craver, Carl F. (2009) "Mechanisms and Natural Kinds", *Philosophical Psychology* 22 (5):575–594.

Devitt, Michael (2008) "Resurrecting Biological Essentialism", *Philosophy of Science* 75 (3):344–382.

Ereshefsky, Marc, and Mohan Matthen (2005) "Taxonomy, Polymorphism, and History: An Introduction to Population Structure Theory", *Philosophy of Science* 72 (1):1–21.

Godfrey-Smith, Peter (forthcoming) "Induction, Samples, and Kinds", in Joseph Keim Campbell, Michael O'Rourke and Matthew H. Slater (eds.), *Carving Nature at its Joints*, Cambridge, MA: MIT Press.

Griffiths, Paul E. (1997) *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.

——— (1999) "Squaring the Circle: Natural Kinds with Historical Essences", in Robert A. Wilson (ed.), *Species: New Interdisciplinary Essays*, Cambridge: MIT Press.

Harman, Gilbert (1965) "Inference to the Best Explanation", *Philosophical Review* 74:88-95.

Kornblith, Hilary (1993) *Inductive Inference and its Natural Ground*. Cambridge: MIT Press.

Lange, Marc (2000) *Natural Laws in Scientific Practice*. New York: Oxford University Press.

——— (2004) "The Autonomy of Functional Biology: a Reply to Rosenberg", *Biology and Philosophy* 19:93–109.

——— (2005) "Laws and Their Stability", *Synthese* 144 (3):415–432.

——— (2009) *Laws and Lawmakers*. Oxford: Oxford University Press.

Lipton, Peter (1996) "Review of Kornblith, *Inductive Inference and its Natural Ground*", *Philosophy and Phenomenological Research* 56 (2):492–494.

Lowe, E. J. (2006) *The Four-Category Ontology: A Metaphysical Foundation for Natural Science*. Oxford: Oxford University Press.

Machamer, Peter, Lindley Darden, and Carl F. Craver (2000) "Thinking About Mechanisms", *Philosophy of Science* 67:1–25.
Magnus, P.D. (MS) "Drakes, Seadevils, and Similarity Fetishism", *Unpublished Manuscript* (2/12/11): Accessed from: http://www.fecundity.com/job/hpcspecies.pdf.
Mayr, Ernst (1963) *Animal Species and Evolution*. Cambridge: Harvard University Press.
Mill, John Stuart (1872) *A System of Logic*. 8th ed. London: Longman, Green, and Co. (Reprinted by Spottiswoode, Ballantyne & Co., 1965).
Millikan, Ruth Garrett (1999) "Historical Kinds and the "Special Sciences"", *Philosophical Studies* 95:45–65.
——— (2000) *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.
Molnar, George (2003) *Powers: A Study in Metaphysics*. Oxford: Oxford University Press.
Norton, John D. (2003) "Causation as Folk Science", *Philosophers' Imprint* 3 (4):1–22.
Okasha, Samir (2002) "Darwinian Metaphysics: Species and the Question of Essentialism", *Synthese* 131:191–213.
Peacocke, Christopher (2004) *The Realm of Reason*. New York: Oxford University Press.
Putnam, Hilary (1975) "Is Semantics Possible?" in his *Mind, Language and Reality: Philosophical Papers*, Cambridge: Cambridge University Press.
Richards, Richard A. (2010) *The Species Problem*. New York: Cambridge University Press.
Ridley, Mark (2003) *Evolution*. 3rd ed. Boston: Blackwell Scientific.
Ruse, Michael (1987) "Biological Species: Natural Kinds, Individuals, or What?", *British Journal for the Philosophy of Science* 38:225–242.
Shriver, Duward F., Peter Atkins, and Cooper H. Langford (1994) *Inorganic Chemistry*. Second ed. New York: W. H. Freeman and Company.
Slater, Matthew H. (2005) "Monism on the One Hand, Pluralism on the Other", *Philosophy of Science* 72 (1):22–42.
Sterelny, Kim (1999) "Species as Ecological Mosaics", in Robert A. Wilson (ed.), *Species: New Interdisciplinary Essays*, Cambridge: MIT Press.
Sterelny, Kim, and James Maclaurin (2008) *What Is Biodiversity?* Chicago: University of Chicago Press.
Strevens, Michael (2008) *Depth*. Cambridge: Harvard University Press.
Whewell, William (1840) *Philosophy of the Inductive Sciences*. London: Parker.
White, Roger (2005) "Explanation as a Guide to Induction", *Philosophers' Imprint* 5 (2):1–29.
Wilson, Robert A. (1999) "Realism, Essence, and Kind: Resuscitating Species Essentialism?", in Robert A. Wilson (ed.), *Species: New Interdisciplinary Essays*, Cambridge: MIT Press.
——— (2005) *Genes and the Agents of Life*. Cambridge: Cambridge University Press.
Wilson, Robert A., Matthew J. Barker, and Ingo Brigandt (2007) "When Traditional Essentialism Fails: Biological Natural Kinds", *Philosophical Topics* 35 (1/2):189–215.
Woodward, Jim (2002) "What Is a Mechanism? A Counterfactual Account", *Philosophy of Science* 69:366–377.