

Acquaintance, Parsimony, and Epiphenomenalism

Forthcoming in S. Coleman, ed., *The Knowledge Argument Then and Now*, Cambridge University Press.

Brie Gertler, University of Virginia

Why does Mary learn something when she leaves the room? One answer, endorsed by some physicalists as well as most dualists, is as follows. Mary learns something because phenomenal knowledge requires direct acquaintance with phenomenal properties. For this reason, there is an epistemic gap between the physical and the phenomenal: phenomenal facts cannot be deduced from physical facts. This is *the acquaintance response* to the Knowledge Argument. The physicalist and dualist versions of the acquaintance response diverge as to whether this epistemic gap reveals an ontological gap between the physical and the phenomenal.

The acquaintance response is, I believe, an especially promising way to make sense of the Mary case. I will not argue for the acquaintance response here, although I do hope to make its appeal clear. My focus will be on teasing out its implications for the debate between physicalism and dualism. Specifically, I will argue that the acquaintance response casts doubt on two claims often made on behalf of physicalism: that physicalism is more parsimonious than dualism, and that no plausible view about mental causation is compatible with dualism.

The paper has three sections. Section 1 explicates the acquaintance response to the Knowledge Argument and outlines Acquaintance Physicalism, the position that combines the acquaintance response with physicalism. In Section 2, I argue that physicalism's claim to greater parsimony is less straightforward than usually assumed, and that the commitments of Acquaintance Physicalism present special obstacles to invoking parsimony in an argument for physicalism. And I show that on an alternative interpretation of parsimony (Sober 2015), physicalism is not more parsimonious than dualism *per se*. Section 3 shows how acknowledging the phenomenon of acquaintance can ease the dualist's problems with mental causation, by dispelling three key objections to epiphenomenalism. The most challenging of these objections is that epiphenomenalism blocks an evolutionary explanation of the fact that events beneficial to the organism are generally pleasurable while harmful events are generally painful. In response, I draw on the relation of acquaintance to describe how pleasures and pains, while themselves epiphenomenal, might nonetheless *explain* positive and negative associations with stimuli. Because these associations affect behavior, they can contribute to fitness. I close by arguing that epiphenomenalism does not threaten human agency.

1. Acquaintance Physicalism

1.1 *The Acquaintance Response to the Knowledge Argument*

The Phenomenal Concept Strategy (PCS) attempts to block the Knowledge Argument's anti-physicalist implications by claiming that, when Mary is released from the room and sees something red, she merely acquires a new way of conceptualizing phenomenal redness. She does not gain knowledge of a previously unknown *property*. The PCS allows that Mary learns something upon her release, but blocks the inference to dualism by claiming that our ways of conceptualizing reality are more fine-grained than the reality we represent.

On one version of the PCS, what Mary gains is simply a new way of *referring to* phenomenal redness, perhaps via an introspective demonstrative such as *this* color (Papineau 2002, Levin 2006). This way of cashing out Mary's new phenomenal concept is what Chalmers calls a "thin" account, as it implies that Mary's new knowledge does not constitute a substantial epistemic advance. (Compare: it is only by visiting the Vatican that I can refer to the Pope's residence by pointing at it. Yet the fact that I can now knowledgeably say "The Pope lives *there*" does not mark a substantial epistemic advance.) The chief worry about thin versions of the PCS is that they cannot accommodate the epistemic intuitions that drive anti-physicalist arguments, such as the intuition that Mary acquires substantial new knowledge when she is released and learns what it's like to see red (Chalmers 2006).

My current concern is not with thin accounts of phenomenal concepts. It is instead with the acquaintance response to the Knowledge Argument, which construes phenomenal concepts as epistemically substantial or "thick". According to the acquaintance response, Mary's new knowledge is a genuine epistemic advance. Phenomenal knowledge—knowledge of what it's like—requires a grasp of the phenomenal property that can be achieved only through direct acquaintance with that property. It therefore requires that the target property is instantiated in one's own experience. This is why Mary gains phenomenal knowledge of redness only by experiencing redness herself; she cannot deduce this from the knowledge she has about redness before her release.

The claim that phenomenal knowledge requires acquaintance explains why there is something that Mary doesn't know about "seeing red" experiences until she is released from her room. Before her release, she can conceptualize phenomenal redness in physical terms—e.g., as *the phenomenal property that ordinary persons experience when their retinas are struck by light with a wavelength of 620-780 nanometers*, and perhaps as *the phenomenal property correlated with neural state N*. But she cannot conceptualize it phenomenally. It is only upon her release, when she is acquainted with an instance of phenomenal redness (seeing a stop sign, perhaps), that she is in a position to conceptualize this property phenomenally.

Generalizing from this case, we get the following principle.

Phenomenal Acquaintance. One conceptualizes a phenomenal property *phenomenally*—i.e., in terms of what it's like—only when one is acquainted with the property, which in turn requires having the relevant phenomenal experience. Phenomenal conceptualization provides for *phenomenal knowledge*: epistemically substantial knowledge of what it's like. Phenomenal knowledge cannot be achieved by deductive inference from non-phenomenal knowledge, e.g., from knowledge involving only physical concepts.

This principle is part of the acquaintance response to the Knowledge Argument. That response is endorsed by most dualists (Nida-Rümelin 1995; Chalmers 2003; Gertler 2012; Goff 2015). It is also accepted by some physicalists (Balog 2012; Howell 2013). I'll call these positions *Acquaintance Dualism* and *Acquaintance Physicalism*, respectively.

1.2 *Acquaintance Physicalism: the view*

Chalmers argues that the acquaintance response to the Knowledge Argument cannot be reconciled with physicalism, since "crucial explanatory elements in the account will not be

physically explainable” (Chalmers 2006, 183). In particular, he thinks that the notion of *acquaintance* resists physical explanation. Levine nicely illuminates why simply invoking acquaintance will not resolve explanatory gap worries. He considers a standard view, on which acquaintance involves an instance of the phenomenal property being contained within the physical structure that realizes the phenomenal concept of that property.

Acquaintance, or cognitive presence, or whatever it is that is supposed to constitute the especially immediate and intimate cognitive relation between phenomenal concepts and their objects, is just that: a *cognitive* relation. It is not at all clear why, or how, *physical* presence [an instantiation of the phenomenal property] translates into *cognitive* presence. (Levine 2006, 162)

The transition from physical containment to awareness—the special kind allegedly afforded by phenomenal concepts—is still an inexplicable transition. It is subject to its own explanatory gap, just as much as is the original relation between phenomenal properties and their physical correlates. (*ibid.*, 163)

Acquaintance Physicalists concede this point but deny that it threatens physicalism. For example, Balog grants that physical theory will not bridge the gap between Mary’s way of conceptualizing phenomenal redness before her release—viz., in physical terms—and the conceptualization made possible by acquaintance with that property. But, she argues, the remaining epistemic gap between physical and phenomenal conceptualizations does not threaten physicalism unless it is assumed that “epistemic gaps always indicate ontological gaps” (Balog 2012, 18). And she thinks physicalists should reject that assumption.

Howell shares this outlook. He acknowledges the force of the Chalmers-Levine worry: “There appear to be truths about acquaintance that cannot be deduced from the complete physical truth” (Howell 2013, 147). And he notes that this non-deducibility is itself an explanatory gap: physical theory cannot explain the relevant facts about acquaintance. But since explanation is itself an epistemic notion, he says, the remaining gap is only epistemic. So it is compatible with physicalism, which is after all an ontological thesis.

The Knowledge Argument uses an epistemic divide—that phenomenal knowledge cannot be achieved by deduction from physical knowledge—to establish an ontological divide between the phenomenal and the physical. Acquaintance Physicalists recognize the epistemic divide but reject the inference to ontological dualism. Acquaintance Physicalism is therefore a version of *a posteriori* physicalism. But whereas most *a posteriori* physicalists assimilate the phenomenal-physical epistemic gap to the epistemic gaps that allegedly characterize other cases of *a posteriori* necessities or identities, Acquaintance Physicalists invoke the distinctive metaphysics of the acquaintance relation to explain what is epistemically special about acquaintance. In acquaintance, a “token of a phenomenal concept applied to current experience is (partly) constituted by that token experience” (Balog 2012, 7). Or “acquaintance imports phenomenal states into propositions” (Howell 2013, 123).

Acquaintance Physicalism contests, and thereby reveals, deep-seated presumptions about the relation between the epistemic and the ontological. These presumptions are operative in Chalmers’ (2006) claim that PCS strategists face a dilemma: *thin* accounts of phenomenal concepts may be explainable in physical terms, but will not do justice to the explanatory gap; *thick* accounts of phenomenal concepts may do justice to the explanatory gap, but will not be

explainable in physical terms. The second horn of this dilemma is premised on the idea that physicalism, an ontological thesis, would be undermined if the epistemic significance of acquaintance were physicalistically inexplicable—an epistemic shortcoming.

Strikingly, these deep-seated presumptions about the relation between the epistemic and the ontological are also present in arguments that favor physicalism. In a recent paper, Veillet presents a dilemma that is in some respects similar to Chalmers' (Veillet 2015). But while Chalmers intends his dilemma to show that physicalists cannot answer the challenge presented by the epistemic gap between the physical and the phenomenal, Veillet intends her dilemma to show that they *need not* answer it.

The first horn of Veillet's dilemma is that knowledge (or the concepts exercised therein) is individuated more finely than its objects (or the properties conceptualized). In that case, Veillet says, the fact that Mary gains new knowledge upon her release poses no threat to physicalism: "the challenge [to physicalism] simply dissolves" (ibid., 2964). The second horn is that knowledge is individuated coarsely. Veillet argues that the only reasonable way to make sense of the idea that phenomenal and physical knowledge differ, *when knowledge is individuated coarsely*, is to say that they are knowledge of different properties. But on this reading of the epistemic gap between physical and phenomenal knowledge, she contends, the claim that there is such a gap begs the question against physicalism. So Mary's new knowledge is either too insubstantial to present a challenge to physicalism, or too substantial to avoid begging the question against physicalism.

Acquaintance Physicalists maintain that knowledge of phenomenal properties is individuated at a finer grain than those properties; in this sense, they endorse the first horn of Veillet's dilemma. But they allow that the Mary case poses a genuine and powerful challenge to physicalism. And while they ultimately deny that the challenge succeeds, they maintain that answering the challenge is a significant task. It ultimately requires commitment to *acquaintance* as an epistemically and metaphysically distinctive phenomenon, to explain why there remains, in Howell's terms, an "epistemological rift" between the physical and the phenomenal. Howell says that by adopting his view,

... we have staved off a metaphysical rift in the world, but that comes at the cost of an epistemological rift. There is a sense in which we cannot fully grasp the physicality of conscious states. It is my view that this is a rather significant admission. (Howell 2013, 173)

Chalmers and Veillet, who disagree on the question of physicalism, agree on how epistemic matters bear on ontological matters in this context. They agree that if Mary's new knowledge is epistemically substantial, then physicalism is in jeopardy. (Their disagreement concerns whether this link between the epistemic and the ontological threatens physicalism or, rather, shows that epistemic arguments such as the Knowledge Argument beg the question.)

Acquaintance Physicalists reject this shared outlook. They claim that Mary's new knowledge is epistemically substantial, and that what explains why phenomenal knowledge cannot be achieved by deduction from physical facts—namely, that some epistemically substantial knowledge requires acquaintance—itself resists physicalistic explanation. So Mary's new knowledge is epistemically substantial enough to pose a real challenge to physicalism. But

physicalism can meet this challenge, by invoking the notion of acquaintance to explain the fundamental epistemic gap between the physical and the phenomenal.

1.3 *Acquaintance Physicalism as an error theory*

The acquaintance response is an appealing way to make sense of the Mary case. As I mentioned above, most dualists accept the acquaintance response. The Acquaintance Dualist construes Mary's epistemic advance as discovering that certain genuine metaphysical possibilities, e.g. that a different phenomenal property correlates with the physical features she associates with *seeing red*, are non-actual.

But Acquaintance Physicalists deny that what Mary learns is that certain metaphysical possibilities are non-actual. They maintain that the possibilities ruled out by her new knowledge—e.g., that the physical features associated with *seeing red* are correlated with phenomenal greenness—are only epistemic or conceptual possibilities, not metaphysical possibilities. On their view phenomenal properties are metaphysically necessitated by (or identical to) physical properties, despite the fact that phenomenal facts cannot be deduced from physical facts. So Acquaintance Physicalism is an error theory. The phenomenon of acquaintance, with its distinctive epistemic and metaphysical characteristics, is invoked to explain why a necessary link between the phenomenal and the physical *seems* contingent.

There is a sort of cognitive block that prevents us, and will always prevent us, from deducing conscious states from physical states. (Howell 2013, 170)

Of course, physicalism would remain puzzling and downright incomprehensible if a perspicuous physicalist explanation of the epistemic gaps themselves was not possible. The crucial element of [my account] is that it provides just such an explanation. It offers the next best thing to a perspicuous explanation of Q [the phenomenal facts] in terms of P [the physical facts], namely, it offers a perspicuous explanation of why we can't have one. (Balog 2012, 20)

Acquaintance Physicalists are driven to an error theory about the apparent contingency of the phenomenal-physical relation by their antecedent commitment to physicalism. This brings us to the motivations for physicalism. Physicalism is thought to possess at least two important advantages over dualism: it is thought to be theoretically preferable to dualism on grounds of parsimony, and it is believed to provide for a more plausible view about mental causation.¹ I address these two alleged advantages in the next two sections, respectively. (*Note to the reader:* Sections 2 and 3 are largely independent. Those interested primarily in the issue of mental causation can move directly to Section 3.)

2. Acquaintance Physicalism and the Question of Parsimony

2.1 *Is physicalism obviously more parsimonious than dualism?*

One factor that seems to favor physicalism is that it is ontologically more parsimonious than dualism. It's not clear how much weight parsimony considerations carry in the current debate. They are sometimes mentioned,² but physicalists more frequently cite concerns about mental causation. Still, it often seems that parsimony concerns are working behind the scenes,

fueling the sense that dualists bear the initial burden of proof because they posit kinds of properties beyond those recognized by the physicalist.

It seems hard to deny that physicalism is more parsimonious than dualism. After all, physicalism is the thesis that all of concrete reality is, or is metaphysically necessitated by, entities that belong to a single kind: the physical. (I will use the term “entity” liberally, to cover properties, events, and things.) Dualism rejects this monism, taking the mental to be an additional basic kind. So physicalism seems clearly more parsimonious: it recognizes fewer basic kinds of entities.³

However, the situation is not as clear as it seems. Physicalism is a genuine monism only if “the physical” names a single kind. But it is far from obvious that there is a principled (non *ad hoc*) way of individuating kinds so that all physical entities belong to a single kind. And if “the physical” comprises numerous fundamental kinds, physicalism may not be significantly more parsimonious than dualism. If the set of “physical” kinds is highly heterogeneous, then dualism may be just as parsimonious as physicalism. For example, accommodating the relevant data in a way that avoids dualism could require positing a *sui generis* basic physical kind. If “the physical” is a heterogeneous motley, then this *sui generis* kind may not be more similar to other physical kinds than a non-physical kind would be. In that case, physicalism and dualism could be equally parsimonious.

2.2 *The epistemic optimism of the “greater parsimony” claim*

Physicalism is standardly regarded as more parsimonious than dualism because “the physical” is taken to name a single kind, or to encompass a relatively homogeneous set of kinds. The claim that physicalism is more parsimonious than dualism seems to me to rest on highly optimistic assumptions about how we conceptualize physical reality. These assumptions fit uneasily with Acquaintance Physicalism’s claim that our way of conceptualizing the phenomenal is fundamentally misleading. I’ll make this case in reference to one formulation of physicalism, though the point applies to at least some others as well.

How should we understand physicalism? Some philosophers argue that the point at issue between physicalists and dualists is whether mentality appears at the fundamental level of ontology (Montero 2001). This may be a reasonable way to interpret the debate. But it does not imply that physicalism is monistic, or even that it is more parsimonious than dualism. The “no fundamental mentality” thesis associated with physicalism simply says that there is a particular kind, the mental, that is not part of fundamental ontology. It says nothing about how many kinds of entity are fundamental. More to the point, that a theory denies that a particular kind is ontologically basic does little to support the claim that it is more parsimonious than a theory that includes that kind in its basic ontology. For as just mentioned, it might be that excluding that kind requires recognizing additional kinds.

A reasonable formulation of physicalism should be linked with our empirically-informed conception of the physical. To deal with Hempel’s dilemma, a formulation of physicalism must be liberal enough to allow that physical theory may advance in unanticipated directions, yet restrictive enough to avoid rendering physicalism trivially true or equivalent to naturalism. The formulation I will adopt begins from the idea that physical theory interprets concrete reality as having a spatiotemporal structure, and individuates entities by their causal powers and relations.

The version of this approach I will use is Howell's. Ignoring some complications, Howell's basic conception of physicalism is as follows (the label "Spatiotemporal-Causal Physicalism" is my own).⁴

Spatiotemporal-Causal Physicalism: [T]he concrete properties and things in this world [metaphysically] supervene upon the properties in this world that are exhausted by their implications for the distribution of things over space and time. (Howell 2013, 53)

Let's call this thesis "STC Physicalism". And let's use the term "STC entities" for the concrete entities it recognizes: that is, for entities that are "exhausted by their implications for the distribution of things over space and time", or metaphysically supervene on properties meeting that description. This conception of physicalism construes phenomenal property dualism as the claim that phenomenal properties are not STC entities. This is a plausible construal. Even those dualists who take the phenomenal to have causal powers will deny that phenomenal properties are causally analyzable, or metaphysically supervene on properties amenable to causal analysis.

How parsimonious is STC Physicalism? This depends on whether "STC entities" names a single kind; or, if not, whether STC entities form a relatively homogeneous set. I don't know how to begin to answer that question. Physicists now suspect that positing novel kinds of fundamental (presumably STC) entities will be needed to solve remaining mysteries, such as the nature of dark matter and how it exerts gravitational force. Arguably, then, no one is currently in a good position to determine whether the basic kinds of STC entities form a homogeneous set.

Of course, the real issue is comparative: whether STC Physicalism is more parsimonious than dualism. The claim that STC Physicalism is more parsimonious than dualism depends on the assumption that the distinction between STC entities and non-STC entities carves nature at a real joint, one that is comparatively deep. This assumption is needed to exclude possibilities like the following. The class of STC entities comprises a heterogeneous mix of dozens of different basic kinds, and the differences between these kinds are just as deep as the difference between STC and (perhaps merely possible) non-STC entities. In that case, STC Physicalism would be only marginally more parsimonious than dualism. And if STC Physicalism required positing a basic physical kind beyond the physical kinds recognized by dualism, then these theories could conceivably be on a par as regards parsimony. So the parsimony question rests on whether the STC / non-STC distinction carves nature at a particularly deep joint.

Let's grant that the sciences generally conceptualize physical entities in terms of their STC features. This approach to conceptualization is informed, it seems, by how we detect physical entities: namely, through their spatiotemporal effects (broadly speaking). It's hard to deny that this way of conceptualizing physical entities has proven amazingly effective in navigating physical reality. STC Physicalism ambitiously claims that this understanding is not only useful, but captures the nature of concrete reality in itself, and does so comprehensively. And the claim that STC Physicalism's ontology is significantly more parsimonious than an ontology that recognizes non-STC entities puts even more stock in our conceptualization practices. For it implies that our way of conceptualizing physical reality, in spatiotemporal-causal terms, reflects an especially deep ontological division, between STC and (possible) non-STC entities.

My purpose is not to deny that physicalism is more parsimonious than dualism. It is only to expose an assumption required to establish that thesis, at least on the STC construal of physicalism. The assumption is that our conceptualization of physical entities, itself derived from our means of observation, reflects genuine and significant divisions in nature. With this assumption in mind, let's return to the acquaintance response to the Knowledge Argument.

2.3 *The epistemic pessimism of Acquaintance Physicalism*

According to the acquaintance response, our phenomenal conceptualization of the property *phenomenal redness* is derived from a particular means of apprehending this property, namely, through direct acquaintance with it. This conceptualization is as a non-STC property: it does not represent phenomenal redness as exhausted by its causal implications for spatiotemporal entities. If it did represent phenomenal redness that way, this conceptualization would be available to Mary before her release. For the same reason, phenomenal facts (facts under a phenomenal conception) are not conceptualized *as* necessitated by physical facts (or facts involving entities at the STC base, viz., those exhausted by their STC implications). To conceptualize Q as necessitated by P is to conceptualize Q in a way that makes it seem deducible from P.

The Acquaintance Dualist allows that the conceptualization afforded by acquaintance captures a truth about phenomenal properties: that they are non-STC entities. By contrast, the Acquaintance Physicalist regards our means of achieving phenomenal knowledge, namely through acquaintance, as fundamentally misleading about the nature of concrete reality.

Acquaintance Physicalism thus rests on two sharply diverging assessments of how we conceptualize concrete reality, as follows. On the one hand, our means of understanding the physical, through detecting effects (and inferences therefrom), shapes our conceptualization of the physical as STC entities. This conceptualization is broadly accurate. And it is comprehensive, in that it encompasses all of concrete reality. On the other hand, our means of understanding the phenomenal, through acquaintance, shapes our conceptualization of phenomenal properties as non-STC properties. This conceptualization is fundamentally mistaken, and coming to terms with it requires an elaborate error theory.

The rationale for these differing verdicts is of course the commitment to physicalism. But if my argument above is correct, one advantage claimed for physicalism—its comparative parsimony—requires a high degree of confidence in our physical conceptualization of concrete reality. Reason to think that STC Physicalism is more parsimonious than dualism rests on the assumption that our conceptualization of concrete reality—in particular, our distinction between STC and non-STC entities—reflects not only a real ontological difference, but one that is especially significant, relative to differences between kinds of STC entities. So an appeal to parsimony puts even more stock in our conceptualization of the physical. An appeal to parsimony would thus increase the already sharp disparity between Acquaintance Physicalism's high confidence in our (perceptually-derived) physical conceptualization of concrete reality, and that view's error theory about our (acquaintance-derived) phenomenal conceptualization of it. Justifying these sharply diverging assessments of our conceptualization practices, in a way that is not *ad hoc*, is a substantial task.

The claim that STC physicalism is more parsimonious than dualism relies on the assumption that our conceptualization practices track genuine, significant divisions in concrete

reality. This is true of parsimony claims concerning physicalism on at least some other formulations as well. To avoid portraying physicalism as trivially true, or equivalent to naturalism, an adequate formulation of physicalism must include some restrictions on “the physical”. These restrictions standardly involve some way of conceptualizing the physical that derives from the basic methods or commitments of physical science. And for the reasons just given, physicalism thus construed has a claim to parsimony only if these conceptualizations track genuine, significant divisions in concrete reality.

Parsimony claims on behalf of physicalism are highly optimistic about one dimension of our conceptualization practices. Acquaintance Physicalism is committed to an error theory about another dimension of our conceptualization practices. The need to provide a principled justification for this disparity constitutes a hurdle for an Acquaintance Physicalist who wishes to invoke parsimony.

2.4 Parsimony and predictive accuracy

Might the case for physicalism’s greater parsimony be made in a way that doesn’t rely on the accuracy of our conceptualization practices? Sober presents such a case, using a model selection framework. But this argument may not help the physicalist. For the framework Sober advances does not associate a theory’s greater parsimony with a greater likelihood of truth; it instead associates parsimony with a theory’s predictive accuracy. Sober expresses skepticism about the association of parsimony with truth, as regards mental ontology.

Placing the mind/body identity theory and dualism within the context of model selection theory requires one to think of the contending theories in terms of their predictive accuracy, not their truth. Metaphysicians may balk at this, proclaiming that they don’t care about predictive accuracy and want only to figure out what is true. I am not arguing against that preference. Rather, my point is that the parsimony argument for the identity theory finds a natural home in the model selection framework. If there is another treatment of the argument that establishes its connection with truth, I do not know what that treatment is. (Sober 2009, 137)

Insofar as our concern is with the truth of physicalism, the results of the model selection framework may not be directly relevant to our discussion.

Still, it’s worth noting that those results do not straightforwardly favor physicalism over dualism. In the model selection framework, models are selected by their Akaike Information Criterion score, a mark of predictive accuracy in which parsimony yields greater predictive accuracy (all else being equal). This framework ties a model’s parsimony to the number of “adjustable parameters” it allows, where this number derives from the number of possibilities compatible with the model. Consider the four logical possibilities represented in a table provided by Sober (2015).

	The brain monitor says “c-fibers are firing.”	The brain monitor says “no c-fibers are firing.”
S pushes the button that says “pain.”	f_1	f_2
S pushes the button that says “no pain.”	f_3	f_4

Versions of physicalism that allow for multiple realizability are compatible with three possibilities— f_1 , f_2 , and f_4 . Of course, these physicalist views do not deny that f_3 is possible; they merely imply that any occurrence of f_3 is due to measuring error. But precisely the same is true of leading versions of dualism, which take the phenomenal to be nomologically linked with the physical. If c-fiber firing nomologically suffices for pain (and these are synchronous), then f_1 , f_2 , and f_4 could occur. And f_3 will not occur unless there is measuring error. So nomological supervenience dualism has precisely the same adjustable parameters as multiple realization physicalism. And so these theories have the same number of adjustable parameters, the measure of parsimony that factors into their predictive accuracy scores.⁵

This result comes about because the model selection framework is insensitive to mere (non-actual) possibilities. It therefore will not distinguish between the metaphysically necessary supervenience associated with versions of physicalism that allow for multiple realizability, and the nomological supervenience associated with leading versions of dualism.⁶ For the same reason, the identity theory’s predictive accuracy score will be shared by a version of dualism according to which c-fiber firing is nomologically necessary as well as nomologically sufficient for pain. These theories are both compatible with f_1 and f_4 , and they will count f_2 and f_3 as due to measuring error. So the identity theory is more parsimonious than a theory allowing multiple realization, and the dualist theory just described is more parsimonious than one on which c-fiber firing is nomologically sufficient but not necessary for pain. The model selection framework enables these comparisons, but it does not imply that physicalism is more parsimonious than dualism *per se*.

2.5 Summing up: Acquaintance Physicalism and the question of parsimony

The parsimony case for physicalism depends on an especially high degree of confidence in our conceptualization practices. Acquaintance Physicalism’s skepticism about our conceptualization of the phenomenal is therefore an obstacle to the parsimony case for physicalism. An alternative approach to parsimony, model selection theory, evaluates parsimony in terms of empirical possibilities recognized by the theory. But because it is not sensitive to merely modal differences, it does not imply that physicalism is more parsimonious than dualism *per se*.

3. Acquaintance and Epiphenomenalism

3.1 Worries about epiphenomenalism

The most frequently cited motivation for physicalism concerns mental causation. The positions compatible with dualism—interactionism, parallelism, and (bottom-up)

epiphenomenalism—are generally thought to be unappealing. I will focus on epiphenomenalism, the position that Jackson defended in his 1982 paper. Philosophers famously express dismay at the prospect of epiphenomenalism. At best it’s “repugnant to good sense” (Lewis 1988, 514); at worst, it’s “the end of the world” (Fodor 1989, 79).

But *why* this prospect should trouble us is unclear. The worry does not seem to be that epiphenomenalism is patently false. The falsity of epiphenomenalism is not an introspective datum. For familiar Humean reasons, no amount of introspective observation could reveal the mental’s causal powers. And Jackson provides a plausible explanation for why phenomenal states seem causally efficacious. So long as physical states with causal potency nomologically suffice for—or, on Jackson’s version, cause—phenomenal states, it will seem that phenomenal states have causal potency. If the firing of c-fibers both nomologically suffices for pain and causes the subject to yell “ouch!”, it might easily seem that the pain causes the yelling. The general point is similar to that just made, about the limitations of model selection theory. Observation of regularities will not distinguish between being (necessitated by or identical to) the cause and merely being nomologically linked with the cause.

In this section, I show how recognizing the role of acquaintance in phenomenal knowledge can mitigate three key objections to epiphenomenalism. The first objection, which others have addressed, is that epiphenomenal states would not be observable. The second objection is that epiphenomenalism amounts to an unpalatable exceptionalism about phenomenal properties or states, and so commitment to epiphenomenalism would compromise the simplicity of a theory. The third objection is, to my mind, the most challenging. It is that epiphenomenalism blocks evolutionary explanations of the fact that events harmful to the organism tend to be painful, while beneficial events tend to be pleasurable. I suggest that acquaintance can help to neutralize this objection as well.

These three objections are broadly epistemic: they concern apparent evidence against epiphenomenalism. I conclude this section by briefly addressing a non-epistemic but serious concern, namely that epiphenomenalism would deprive us of agency.

3.2 Epiphenomenalism and Knowledge of Phenomenal Properties

Watkins worries that “on Jackson’s epiphenomenalism we cannot know of our own qualitative experiences If qualia are not causally efficacious, then my beliefs and memories would be just as they are whether there were qualia or not.” (Watkins 1989, 160). This is known as the *self-stultification* objection. It challenges the Knowledge Argument itself, since it suggests that if epiphenomenalism were true then Mary couldn’t know what it’s like to see red even after her release (Campbell 2003).

The acquaintance thesis provides for a response to the self-stultification objection (Nagasawa 2010), since knowledge by acquaintance is achieved through a non-causal route. As Balog says, in acquaintance an instance of the property serves as a mode of presentation of the property itself. So the relation between the thought about the property—“*this* is what it’s like to see red”—and the property instance is direct and non-causal. This constitution relation enables the subject to directly (non-causally) grasp the experience’s phenomenal character. The experience needn’t cause the awareness of it, since it partly constitutes that awareness. When one grasps a phenomenal property by acquaintance, one’s judgment regarding that experience is tied to reality directly: an instance of the phenomenal property is part of the judgment, which

expresses how that reality epistemically seems to the subject. (For details and defense of this claim, see Gertler 2012.)

De Brigard expresses a worry related to the self-stultification objection. He worries that epiphenomenalists cannot explain “how the content of a mental state ... can refer to a particular sensation” (De Brigard 2014, 125). He reasons that such reference must be causal or, at the very least, nomological. This worry can also be assuaged by the claim that we can grasp phenomenal properties through acquaintance. In knowledge by acquaintance, the phenomenal properties serve as the mode of presentation for their referents. Because reference is grounded in a *constitution* relation, it is more direct and epistemically secure than a causal relation.

Another worry, also related to the self-stultification objection, concerns knowledge of other minds. Jackson’s response to this worry implicitly relies on the assumption that I can be aware of my own phenomenal states. He argues that one can know that someone else is in pain by inferring from a behavioral effect back to its cause (a neural event, say), and then inferring the occurrence of another standard effect of that neural event (a phenomenal state). For example, if I hear you yell “ouch!”, I can infer that you are undergoing a neural event of the type that standardly causes pain. Jackson does not explain how the association between the behavior and the phenomenal state is originally established, but presumably this can be established only by using my own phenomenal states; I know that I yell “ouch!” only when I feel pain. So this response to the worry about other minds implicitly assumes the capacity for self-knowledge. Fortunately for this account, one can make this association without accepting epiphenomenalism. One may mistakenly believe that the association between yelling “ouch!” and pain is between an effect and its cause, when in fact it’s between two effects of a single cause.

3.3 *Epiphenomenalism and simplicity*

Papineau suggests that accepting epiphenomenalism would violate simplicity considerations.

If epiphenomenalism [about conscious states] were true, then the relation between mind and brain would be like nothing else in nature. After all, science recognizes no other examples of “causal danglers”, ontologically independent states with causes but no effects. So, given the choice between epiphenomenalism and materialism, standing principles of scientific theory choice would seem to favour materialism. (Papineau 2002, 23)

This argument relies on an inference, from the fact that science doesn’t recognize non-mental epiphenomena, to the conclusion that there aren’t any. But that inference is unjustified. After all, standard scientific methods detect entities by their effects. Indeed, it is the idea that all events are detected by their effects that motivates the previous worry about self-knowledge. So standard scientific methods couldn’t detect epiphenomena even if they were present. In such cases, where the evidence is insensitive to the phenomenon in question, Occam’s Razor does not prescribe denying that the phenomenon occurs. To use Sober’s terms, what applies in this case is not the “razor of denial” but instead the “razor of silence” (Sober 2015).

3.4 *Epiphenomenalism and Natural Selection*

To my mind, the most serious objection to epiphenomenalism is the explanatory challenge raised by William James (among others). If phenomenal properties are epiphenomenal, then the “match” between a condition’s consequences for the organism, and the phenomenal properties associated with it, would seem to be a matter of mere coincidence.

[I]f pleasures and pains have no efficacy, one does not see... why the most noxious acts, such as burning, might not give thrills of delight, and the most necessary ones, such as breathing, cause agony. (James 1890, p. 144)

Robinson aptly refers to this fact, that beneficial events are generally pleasurable while harmful events are generally painful, as “the hedonic/utility match” (Robinson 2007). The natural explanation of this match is an evolutionary one: a creature that feels pain on encountering harmful stimuli will be more likely to withdraw and, hence, more likely to survive long enough to reproduce. But if phenomenal properties are not causally efficacious—e.g., if pain cannot cause a creature to withdraw from a harmful stimulus—then the fact that an experience’s hedonic value matches its utility, for a particular creature, will not affect that creature’s chance of surviving to reproduce. So epiphenomenalism blocks this natural explanation of the hedonic/utility match.

Jackson (1982) responds to a different objection concerning evolution. This is the objection that epiphenomenal properties, of any sort, would not be present in a species that evolved through natural selection. Jackson argues that epiphenomenal properties could be a byproduct of something that *is* selected for, just as the heaviness of a polar bear’s coat is a byproduct of the coat’s warmth. This response shows that the presence of epiphenomenal properties is compatible with evolution. But it doesn’t explain the hedonic/utility match.

In fact, the idea that epiphenomenal properties are merely byproducts of something that is selected for strengthens James’ objection. Jackson uses the example of the polar bear’s coat because, from a fitness perspective, having a heavy coat is disadvantageous; hence, it is not a property that would be selected for. This disadvantage is presumably outweighed (as it were) by the more significant benefit of added warmth. The argument that epiphenomenal qualia may be mere byproducts is intended to explain why they could be present, consistent with not being selected for. If the polar bear’s heavy coat can yield to an evolutionary explanation despite this property’s having a negative effect on fitness, surely qualia can yield to an evolutionary explanation despite having *no* effect on fitness. But if qualia have no effect on fitness, then the hedonic/utility match would appear to be a sheer coincidence.⁷

Adequately dealing with James’ objection requires showing how the hedonic/utility match can improve fitness, and hence be selected for, even if phenomenal properties are epiphenomenal. I propose that the fact that an experience has a certain hedonic value can help to explain a creature’s behavior even if the phenomenal properties constituting the experience are causally inert. My proposal, which is admittedly speculative, rests on the idea that the subject’s acquaintance with causally inert phenomenal states can give rise to attitudes or associations with causal efficacy.

Let’s begin with a mundane case in which the hedonic value of an experience helps to explain what someone does.

Alexis tries cilantro for the first time and discovers that, alas, she belongs to the sizeable minority of people for whom cilantro tastes like soap. She says “yuck!” and vows never to eat cilantro again.

It’s tempting to say that the phenomenology involved—the yuckiness—*causes* Alexis to say “yuck!” and to subsequently avoid eating cilantro. But epiphenomenalism is incompatible with this claim. This is the nub of the current puzzle. Epiphenomenalism blocks a natural explanation of our avoidance of unpleasant stimuli, and thereby blocks a natural explanation of why the hedonic/utility match increases the odds that a creature will survive long enough to reproduce.

Here is my proposal. Acquaintance with the yucky taste leads Alexis to realize that cilantro tastes terrible. Her judgment to this effect incorporates an instance of the yucky phenomenal property, which serves as the property’s mode of presentation.⁸ It is structured as follows.

Eating cilantro causes this <yucky gustatory phenomenology> sensation, which is unpleasant.

The gustatory phenomenology does not *cause* Alexis to judge that this sensation is unpleasant. She registers its unpleasantness directly, through her (non-causal) acquaintance with the phenomenology. So the phenomenology—the yuckiness—can help to explain why she judges that the taste is unpleasant, without causing that judgment.

With the judgment *eating cilantro causes this unpleasant sensation*, Alexis associates eating cilantro with a phenomenal property, roughly *this* (yucky) *taste*. But she also classifies it in a more general way, as *unpleasant*. This more general association abstracts away from the specific way in which eating cilantro is unpleasant—that is, its particular phenomenology. It is a relatively generic negative association, signaling that eating cilantro is to be avoided.

I suggest that this negative association is causally responsible for Alexis’s decision to avoid cilantro and her future cilantro-avoidance behavior. The gustatory phenomenology helps to explain why Alexis judges that eating cilantro is unpleasant, and hence to explain her negative association with eating cilantro. But as just outlined, the phenomenology’s contribution to explaining these things is non-causal. And what causally contributes to Alexis’s cilantro-avoidance behavior is only the negative association, not the phenomenology that explains it.

In order for this proposal to remain compatible with epiphenomenalism, the negative association cannot owe its causal efficacy to phenomenology. And this seems plausible. After all, associations often operate unconsciously, as simple behavioral dispositions; and such dispositions can persist beyond any memory of their basis. Alexis need not be aware of the reason she avoids cilantro. When asked why she does this, she may be unable to say more than “I don’t like it” – where this is not tied to any phenomenological memory. In fact, Alexis’s belief that she dislikes cilantro may itself be inferred from her disposition to avoid it. The habit of avoiding cilantro may be so ingrained that she supposes, reasonably enough, that she must have once tried it and disliked it.

But negative associations, when triggered, often do involve phenomenology. For example, when Alexis sees a cilantro-flavored dish listed on a restaurant menu, she may remember the yucky taste. The presence of that memory can easily make it seem, to Alexis, that

the phenomenology causally contributes to her avoidance of cilantro. But on my proposal this is a mistake. The phenomenological memory may contribute to non-causally explaining her aversion, just as the gustatory phenomenology helps to explain her judgment in the original case. But what causes Alexis to avoid cilantro is only the negative association with cilantro, an association that abstracts away from the phenomenological details.

So I propose that an experience's hedonic value can contribute to explaining the subject's behavior, compatibly with epiphenomenalism. The subject can register an experience's hedonic value by acquaintance with the phenomenal properties involved. In registering an experience's hedonic value, the subject associates the relevant stimulus (such as eating cilantro) with a hedonic value (in this case, negative). Crucially, the phenomenal properties of the experience do not *cause* the association with a hedonic value: rather, those properties are *experienced as* having a certain value, and so in grasping those properties through acquaintance the stimulus is associated with that value. This association, which may initially take the form of a judgment, is causally efficacious. It causally contributes to current behavior, e.g., deciding never to eat cilantro again. And it generates a disposition to avoid cilantro in the future.

In this way, acquaintance provides a bridge from the phenomenal properties of experience—instances of which partly constitute judgments about them—to causally efficacious associations or attitudes that are not phenomenological (or, at least, do not owe their causal powers to their phenomenology).⁹

Now to the heart of James' challenge. The hedonic/utility match, in humans, would be a remarkable coincidence unless this match contributed to fitness in our more primitive ancestors. Consider a primitive creature with the kind of hedonic/utility match we've been discussing. Let's call this creature Match. Match experiences pleasure when (to use other examples from James) filling an empty stomach or resting after fatigue, and pain when undergoing starvation or incurring tissue damage.

Suppose that Match moves close to a flame and experiences a painful burning sensation. Crucially, the phenomenal quality of the burning sensation does not *cause* it to be experienced as painful; it *constitutes* that way of experiencing it. The phenomenology thereby helps to explain, but does not cause, Match's association of proximity to flame with *unpleasant* or *to be avoided*. This negative association abstracts away from the specific kind of unpleasant phenomenology linked with proximity to flame—that it's a burning sensation rather than, say, a stabbing sensation. And it causes Match to avoid flames in the future: Match maintains a safe distance from flames. Seeing a flame may trigger a memory of the burning sensation. But just as with Alexis' memory of the yucky taste, the phenomenology involved in this memory does not cause Match's avoidance behavior. At most it contributes (non-causally) to explaining this behavior, in the way outlined earlier.

This case illustrates how an experience's hedonic value can contribute to explaining a primitive creature's behavior, compatibly with epiphenomenalism. The creature registers the event's hedonic value by experiencing it *as* pleasant (positive) or unpleasant (negative). The creature associates the relevant stimulus with this hedonic value. Crucially, the phenomenology of the experience does not *cause* the association with a hedonic value. Rather, the creature registers the experience *as* pleasant or unpleasant directly, through acquaintance with the relevant phenomenal properties. Associating a stimulus with phenomenal properties *experienced as*

having a certain hedonic value thus means associating the stimulus with that value. This association can causally contribute to current behavior, e.g. Match's withdrawing from the flame. And it can generate dispositions, such as the disposition to avoid proximity to flames in the future.

The explanation of the hedonic/utility match is then straightforward. The negative association with flames causally contributes to Match's avoiding close proximity to flames, and thereby to Match's survival. Things are different for Mismatch, a primitive creature in which the hedonic/utility relation is reversed. Mismatch experiences a pleasant sensation when close to flames. Mismatch thus develops a positive association with proximity to flame, and does not avoid or withdraw from them. As compared with Match, Mismatch will be less likely to survive long enough to reproduce. (But we needn't feel too badly: Mismatch presumably dies a blissful death.)

3.5 *Does epiphenomenalism deprive us of agency?*

I have argued that recognizing the role of acquaintance in grasping phenomenal properties undercuts the epistemic arguments against epiphenomenalism. But I doubt that it is epistemic considerations that cause such dismay among philosophers like Fodor and Lewis. It's instructive to consider the context of Fodor's "end of the world" lament.

[I]f it isn't literally true that my wanting is causally responsible for my reaching, and my itching is causally responsible for my scratching, and my believing is causally responsible for my saying, ... [ellipsis in original] if none of that is literally true, then practically everything I believe about anything is false and it's the end of the world. (Fodor 1989, p. 79)

Fodor worries that epiphenomenalism would overturn some of his deeply held and cherished beliefs; and the beliefs in jeopardy concern *action*. I suspect that his dismay stems from the idea that, if epiphenomenalism is true, then in a real sense we lack agency.

The relationship between agency and epiphenomenalism is a weighty topic, and I cannot hope to do it justice here. But I will make a few brief points.

First, because the Knowledge Argument aims to establish dualism only about the phenomenal, the epiphenomenalism at issue concerns only the phenomenal. But some of Fodor's worries concern the causal efficacy of states that, on most accounts, are not phenomenal. Even if desires and beliefs have associated and proprietary phenomenology, on most accounts they are not exhausted by their phenomenology. So long as my desires, beliefs, and intentions do not owe their causal powers to their phenomenal character, epiphenomenalism about the phenomenal does not threaten most ordinary agency. It does not threaten the idea that my desire for water, and my belief that the cup contains water, causally contribute to my reaching for the cup.

Fodor also mentions itching. Epiphenomenalism implies that the phenomenal character of the itch cannot cause the scratching. However, the kind of proposal sketched in the previous section might explain how, consistently with epiphenomenalism, "I scratched because of an itch" could be true. I grasp the feeling through acquaintance, and think "*this <itchy> feeling would be relieved by scratching*"; that thought is a way of thinking "I have a feeling that would be relieved by scratching". That description is very intellectualized, and so dubious—especially

considering that my dog seems to scratch when feeling an itch. But as before, a primitive analogue is available. In the primitive case, the sensation is brutally associated (either innately or through experience) with relief that comes from scratching or just the urge to scratch. So as with the previous cases, a non-causal relation of acquaintance provides a bridge from the *itching* sensation to causally efficacious attitudes or associations that do not owe their causal efficacy to phenomenal character.

Even putting aside that proposal, I confess that I am not very troubled by the idea that some of my scratching is not caused by my itching. What I have in mind is scratching that is *unreflective*, unmediated by deliberation—e.g. about whether scratching might inflame a mosquito bite. Suppose that unreflective scratching is a non-agential reflex directly triggered by the neural basis of itching, rather than by (the feel of) the itch itself. The comparison with my dog suggests, to me at least, that that possibility isn't cause for deep dismay.

What would be dismaying is the idea that we lack a deeper kind of agency, the kind associated with the capacity to act on intentions formed through reflective deliberation about our values and goals. Most contemporary accounts of agency, both compatibilist and libertarian, locate our agency in a capacity of that sort. But the causal factors at issue in those accounts are not phenomenal states, or at least are not states exhausted by phenomenal character. They are largely attitudes: beliefs, desires, (attitudes towards) values, intentions, hopes, etc. Deliberation may involve conscious attitudes of these types, and conscious attitudes may have associated—even proprietary—phenomenology. But since the epiphenomenalism under consideration is only epiphenomenalism about the phenomenal, it does not imply that conscious attitudes are causally inefficacious. It implies only that such attitudes derive their causal powers from something other than their phenomenal character.

Even those who emphasize the importance of consciousness to agency generally deny that it is *phenomenal* consciousness that matters. For example, in arguing that responsibility for action requires consciousness of one's reasons, Levy makes it clear that what he has in mind is *access* consciousness rather than *phenomenal* consciousness. "What is at issue is the availability of certain kinds of representations to the agent, not whether they experience qualia" (Levy 2013, 213). And the issues raised by the Libet cases and others, about whether the time lag of conscious intentions threatens agency, also concern access consciousness—that is, consciousness *of* intentions (Holton 2004).

3.5 Summing up: how acquaintance quiets worries about epiphenomenalism

I have argued that invoking acquaintance can defuse apparent evidence against epiphenomenalism. Those dismayed by the prospect of epiphenomenalism may find solace in my proposal that acquaintance can forge a bridge between phenomenal properties and causally efficacious attitudes. Even if my itching doesn't cause my scratching, it may well be that my itching *explains* my scratching. And epiphenomenalism poses no threat to standard models of agency, which require at most that my decisions are the product of values of which I am aware, or are formed on the basis of accessible reasons.

Conclusion

Many philosophers share the intuition that Mary makes a substantial epistemic advance the first time she sees something she knows to be red. The acquaintance response is an appealing way to make sense of this intuition. I have suggested that the acquaintance response weakens the case for physicalism in two ways. First, the acquaintance response commits physicalists to an error theory about how we conceptualize the phenomenal. This commitment fits uneasily with the robust confidence in our conceptualization practices required to justify the claim that physicalism is more parsimonious than dualism. Second, the idea that we can grasp phenomenal properties through acquaintance goes a long way in neutralizing concerns about epiphenomenalism. It thereby eases the worry that dualism is not compatible with any plausible view about mental causation. This worry is further eased by coming to appreciate that epiphenomenalism does not threaten standard models of agency.

That some physicalists have recently embraced the acquaintance response is strong evidence of its appeal. If my arguments here succeed, the acquaintance response lessens the appeal of physicalism in ways not previously recognized.¹⁰

Works cited

Alter, Torin & Nagasawa, Yujin, eds., (2015). *Consciousness in the Physical World: Perspectives on Russellian Monism*. Oxford University Press.

Alter, Torin & Walter, Sven, eds., (2006) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press.

Balog, Katalin (2012). In Defense of the Phenomenal Concept Strategy. *Philosophy and Phenomenological Research* 84: 1-23.

Campbell, Neil (2003). An inconsistency in the knowledge argument. *Erkenntnis* 58: 261-266.

Chalmers, David J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

--- (2003). The content and epistemology of phenomenal belief. In Quentin Smith & Aleksandar Jokic, eds., *Consciousness: New Philosophical Perspectives*. Oxford University Press 220-72.

--- (2006). Phenomenal concepts and the explanatory gap. In Alter & Walter.

Corabi, Joseph (2008). Pleasure's role in evolution: A response to Robinson. *Journal of Consciousness Studies* 15: 78-86.

De Brigard, Felipe (2014). In defence of the self-stultification objection. *Journal of Consciousness Studies* 21: 120-130.

Fodor, Jerry (1989). Making mind matter more. *Philosophical Topics* 17: 59-79.

Gertler, Brie (2012). Renewed Acquaintance. In Declan Smithies & Daniel Stoljar (eds.), *Introspection and Consciousness*. Oxford University Press 89-123.

Goff, Philip (forthcoming). Real acquaintance and physicalism. In Paul Coates & Sam Coleman, eds., *Phenomenal Qualities: Sense, Perception and Consciousness*. Oxford University Press.

- Holton, Richard (2004). Review of Daniel Wegner, *The illusion of conscious will*. *Mind* 113: 218-221.
- Howell, Robert J. (2013). *Consciousness and the Limits of Objectivity: The Case for Subjective Physicalism*. Oxford University Press.
- James, William (1890). *The Principles of Psychology*. Dover Publications.
- Levin, Janet (2006). What is a phenomenal concept? In Alter & Walter.
- Levine, Joseph (2006). Phenomenal concepts and the materialist constraint. In Alter & Walter.
- Levy, Neil (2013). The Importance of Awareness. *Australasian Journal of Philosophy* 91: 221-229.
- Lewis, David (1988). What experience teaches. Reprinted in Lycan, ed., *Mind and Cognition*, Oxford: Blackwell, 499–519.
- Mele, Alfred (2013). Unconscious decisions and free will. *Philosophical Psychology* 26: 777-789.
- Melnyk, Andrew (2003). *A Physicalist Manifesto: Thoroughly Modern Materialism*. Cambridge University Press.
- Montero, Barbara (2001). Post-physicalism. *Journal of Consciousness Studies* 8: 61-80.
- Fodor, J. (1989). Making mind matter more. *Philosophical Topics* 17: 59-79.
- Nida-Rümelin, Martine (1998). On belief about experiences: An epistemological distinction applied to the knowledge argument against physicalism. *Philosophy and Phenomenological Research* 58: 51-73.
- Papineau, David (2002). *Thinking About Consciousness*. Oxford University Press.
- Pitt, David (2004). The phenomenology of cognition, or, what is it like to think that P? *Philosophy and Phenomenological Research* 69: 1-36.
- Polger, Thomas W. (2011). Are sensations still brain processes? *Philosophical Psychology* 24: 1-21.
- Robinson, William (2007). Evolution and epiphenomenalism. *Journal of Consciousness Studies* 14: 27-42.
- Sober, Elliott (2009). Parsimony Arguments in Science and Philosophy—A Test Case for Naturalism. *Proceedings and Addresses of the American Philosophical Association* 83: 117-155.
- Sober, Elliott (2015). *Ockham's Razors: A User's Manual*. Cambridge University Press.
- Veillet, Bénédicte (2015). The cognitive significance of phenomenal knowledge. *Philosophical Studies* 172: 2955-2974.
- Watkins, Michael (1989). The knowledge argument against the knowledge argument. *Analysis* 49: 158-60.

¹ These are general considerations that drive the preference for physicalism. Since Acquaintance Physicalism is a position about how to reconcile physicalism with the relevant anti-physicalist intuitions, Acquaintance Physicalism is neutral as to the various arguments for physicalism. Howell rests his case for physicalism largely on concerns about mental causation.

² “The identity theory should be favored for broadly theoretical reasons. Whether simplicity and parsimony are the best reasons to favor the theory is another matter. They may be. But they are at

least examples of such reasons” (Polger 2011, 15). Papineau (2002) and Melnyk (2003) also appeal to parsimony in their arguments for physicalism.

³ Dualism incurs additional complexity in explaining the relations between these two basic kinds, e.g., by positing psychophysical laws in addition to the physical laws accepted by both physicalists and dualists (Chalmers 1996). Now if physical entities are individuated by their causal features, then the number of laws may not vary independently of the number of kinds of entities. So the number and complexity of laws will be connected to the issues raised in this section, concerning the degree of heterogeneity among physical entities. In any case, my discussion here is limited to the question of ontological parsimony, namely, the number of basic kinds recognized by a theory.

⁴ Howell later qualifies this formulation, to allow categorical properties not exhausted by causal or dispositional implications. On his view, such properties are *physical* only if they are *individuated by* their dispositional implications, relative to a world (Howell 2013, pp. 25-30).

⁵ Since each allows for three possibilities, they each have two adjustable parameters: the number of adjustable parameters corresponds to the variables needed to fix the remaining probabilities, once an initial probability has been fixed.

⁶ Sober makes a similar point, about nomological supervenience preventing the generation of AIC scores, in another context (Sober, 262-63, esp. note 18). Sober says that on the model selection framework, dualism is less parsimonious than either the identity theory or physicalist functionalism. But the version of dualism he considers allows phenomenal properties to vary independently of physical properties: that version of dualism thus has an additional adjustable parameter, corresponding to f_3 .

⁷ Robinson (2007) argues that an evolutionary explanation of the hedonic/utility match need only make sense of the idea that sensations with hedonic value, such as pleasure, could be correlated with certain kinds of behavior (e.g., pursuit of certain stimuli). And this is accomplished with the hypothesis that pleasure is causally linked with neural events that produce the behavioral effects in question. But as Corabi (2008) observes, Robinson’s proposal leaves a crucial puzzle unresolved: why it would be more likely for events benefitting the organism to cause pleasure than to cause pain.

⁸ In other words, how the experience seems to Alexis *epistemically* is constituted by its qualitative nature, that is, its phenomenal reality. This is not to say that how the property epistemically seems to Alexis *exhausts* its phenomenal reality. As the famous case of the speckled hen demonstrates, phenomenal appearances outstrip epistemic appearances. (I discuss these issues, and the metaphysics of the acquaintance relation, in Gertler 2012.)

⁹ This claim, that Alexis’s negative association with eating cilantro is causally efficacious, faces an objection. The association at issue seems initially to be a kind of judgment, to the effect that *eating cilantro causes something unpleasant*. (This judgment presumably yields a dispositional belief with this content and a disposition to avoid cilantro.) Some have argued that judgments have phenomenal character. And given the assumption of epiphenomenalism, Alexis’s negative association with eating cilantro cannot owe its causal powers to phenomenology. However, epiphenomenalism implies only that any phenomenology tied to the judgment is irrelevant to its causal powers. So long as judgments do not owe their efficacy to their causal powers, they can both have phenomenal character—even what Pitt calls “proprietary phenomenology”—and be causally efficacious, compatibly with epiphenomenalism.

Still, my proposal is incompatible with some stronger claims on behalf of cognitive phenomenology. E.g., if the judgments or associations I describe have their causal powers in virtue of their intentional contents, then (on the assumption of epiphenomenalism) intentional content cannot be exhaustively constituted by phenomenology. (I return to this issue in the next subsection.)

¹⁰ For valuable comments on drafts of this paper, I'm indebted to Robert Howell, Walter Ott, and especially Sam Coleman.