

The Value of Consciousness: A Propaedeutic

Uriah Kriegel

While we rehearse our measureless wealth, it is for thee, dear Mother!
We own it all and several to-day indissoluble in Thee;
—Think not our chant, our show, merely for products gross, or lucre
—it is for Thee, **the Soul**, electric, spiritual!
Our farms, inventions, crops, we own in Thee! cities and States in Thee!
Our freedom all in Thee! our very lives in Thee!

Walt Whitman, "Song of the Exposition"

Introduction/Abstract

Recent work within such disparate research areas as the epistemology of perception, theories of well-being, applied ethics (especially medical ethics and animal ethics), the philosophy of consciousness, and theories of understanding in philosophy of science and epistemology has featured disconnected discussions of a single underlying question: What is the value of consciousness? The purpose of this paper is to construct a theoretical framework in which this question could be addressed systematically and clearly, as well as to explore and review possible elements of an answer.

PART I. SOME QUESTIONS

What is the value of consciousness? In some, the very question is liable to conjure a warm and fuzzy feeling of meaningfulness and depth; in others, it might inspire a vague

suspicion of irrational attachment to mystical or supernatural themes. But what does the question exactly *mean*? My goal in this first part of the paper is to ‘factorize’ the intimidatingly vague and profound-sounding question ‘What is the value of consciousness?’ into a substantial number of much more precise and more tractable sub-questions. I proceed by first highlighting certain ambiguities and uncertainties attending the relevant notions of (i) value, (ii) consciousness, and even (iii) having, then using corresponding disambiguations and precisifications to formulate questions about the value of consciousness more clearly amenable to theoretical treatment.

1. Three Kinds of Value: Epistemic, Ethical, and Aesthetic

There is a classical conception of philosophy that casts it as attempting to explain in a unified manner the true, the good, and the beautiful. The conception is familiar from medieval philosophy, with its theme of *unum, verum, bonum, pulchrum*. But even in traditional Hindu philosophy, for instance, ‘the supreme one’ is commonly characterized as *Satyam Shivam Sundaram* – truth, goodness, and beauty.

This formulation betrays an implicit commitment to monism about three kinds of value: truth monism about epistemic value, goodness monism about ethical value, and beauty monism about aesthetic value. Since none of these is *analytically* true, a more judicious – if less pithy! – formulation would be ‘the epistemically valuable, the ethically valuable, and the aesthetically valuable.’

What I just called ‘truth monism about epistemic value’ is none other than the view often called in epistemology ‘veritism’ (Goldman 2001 Ch.3). We may formulate it as follows:

(TMEV) Truth is the only intrinsic epistemic value; all other epistemic value is merely instrumental.

The idea is that having true beliefs is the only thing epistemically valuable in and of itself. All other epistemic achievements are epistemically valuable only insofar as they are instrumental in bringing about true beliefs. This sort of veritism is widely defended in contemporary epistemology (especially by reliabilists, but not only). Still, other

intrinsic epistemic values have occasionally been proffered, notably *understanding* (Kvanvig 2003). The idea is that understanding something is not analyzable in terms of having true beliefs, and even if understanding something is instrumental in bringing about true beliefs, the value of understanding is not *exhausted* by the fact that it is. Now, one might propose understanding (or some third epistemic good) as supplanting truth in an alternative version of monism about epistemic value; typically, though, understanding is cited as an element in a *pluralist* account of epistemic value. (Note well: I use the expression 'intrinsic' value to denote non-instrumental value, that is, value that does not accrue to something in virtue of its bringing about something else of value; I do *not* use it to denote a value that something has independently of any relations to other items. Another term often used for this is 'final value.')

Corresponding to the above formulation of veritism we may offer a parallel formulation of 'goodness monism about ethical value':

(GMEV) Goodness is the only intrinsic ethical value; all other ethical value is merely instrumental.

Here it is certainly possible to hear the claim as analytically true. But it is not mandatory. Thus, debates on the relationship between the right and the good (starting with Sidgwick 1874) are naturally taken to have implications for GMEV. Many philosophers – especially consequentialists – hold that an action is right just if the state of affairs it leads to, or is intended to lead to, is good. But someone who argues that an action's being right is not analyzable in terms of the goodness of its (intended) consequences may well go on to claim that the action's rightness is ethically valuable in and of itself, thus embracing pluralism about intrinsic ethical value.

Of the three monisms presupposed in 'the true, the good, and the beautiful,' beauty monism about aesthetic value is by far the most contentious (see Danto 2002 for review). Indeed, its rejection has become something of an orthodoxy in contemporary aesthetics. The view may be formulated as follows:

(BMAV) Beauty is the only intrinsic aesthetic value; all other aesthetic value is merely instrumental.

The idea here is that a novel's being delightfully perceptive, a landscape's being sublime, and a painting's being compelling are all aesthetically valuable only insofar as they contribute (whether causally or constitutively) to the novel's, landscape's, or painting's beauty. This is vehemently rejected by a majority of contemporary philosophers of art; a bountiful pluralism of intrinsic aesthetic values reigns instead.

The choice between monism and pluralism in each of these areas is not what concerns me here. What I want to highlight is just that they are three distinct areas – they concern different *kinds* of value. These three values are fundamentally different and incommensurable: one may find a certain quantity Q of ethical value more valuable than some quantity Q^* of epistemic value, but this could only mean that Q is more ethically or more aesthetically valuable than Q^* ; no meaning attaches to the question as to which quantity is more valuable *simpliciter*.

The distinctions between (a) epistemic, ethical, and aesthetic value and (b) intrinsic and instrumental value means that the question 'What is the value of consciousness?' folds within it the following six distinct questions:

- What is the intrinsic epistemic value of consciousness?
- What is the intrinsic ethical value of consciousness?
- What is the intrinsic aesthetic value of consciousness?
- What is the instrumental epistemic value of consciousness?
- What is the instrumental ethical value of consciousness?
- What is the instrumental aesthetic value of consciousness?

Thus we can already factorize our original question into six separate and more precise questions.

In fact, however, the sub-questions about ethical value can be further factorized. At a most basic level, we engage in ethics to answer the following simple question: How should I live? To answer this question, it is useful to address separately three issues about the good life:

- 1) What makes a life good *for the one who lives it*?
- 2) What makes a life *morally* good?

- 3) What is the relationship between a morally good life and a life good for the one who lives it?

Let us say that that which makes life good for the one who lives it has *prudential value*, while that which makes life morally good has *moral value*. Then we should really distinguish *eight* sub-questions about the value of consciousness: concerning epistemic, prudential, moral, and aesthetic value – intrinsic or instrumental.

It is important to appreciate that questions about intrinsic value call for a different kind of answer from questions about instrumental value. To answer the question ‘What is the intrinsic ethical value of consciousness?’, for instance, we would have to say which of the intrinsic ethical values is had by consciousness (where ‘none’ counts as an answer). In contrast, to answer the question ‘What is the instrumental ethical value of consciousness?’, we would have to say in what way consciousness is instrumental to the promotion of (which) intrinsic ethical value(s).

2. Varieties of Consciousness

Like the term ‘value,’ the term ‘consciousness’ introduces potential ambiguities. Most famously, Ned Block (1995) warned against conflating ‘phenomenal’ and ‘access’ consciousness. The basic idea is that a mental state can be said to be conscious, in one sense, because it is something the subject experiences (there is something it is like for her to be in that state), and in another sense, because it is functionally highly integrated and influences information processing across many parts of the cognitive system. It is possible, of course, that the two notions may be coextensive, or perhaps even pick out the same property. All the same, they are different notions and therefore can be used to pose different questions.

Given this, each of the eight questions distinguished in §1 could itself be factorized into two sub-questions. For instance, the question ‘What is the instrumental epistemic value of consciousness?’ should give way to:

- What is the instrumental epistemic value of *phenomenal* consciousness?
- What is the instrumental epistemic value of *access* consciousness?

We are thus led to *sixteen* distinct questions about the value of consciousness (and perhaps more, in case there are other notions of consciousness besides phenomenal and access consciousness).

It is widely thought that phenomenal consciousness is a generic property that admits of several species. Perceptual phenomenal consciousness is one species (which incidentally divides into several *subspecies*: visual, auditory, etc.). Algedonic phenomenal consciousness is another species (itself dividing into pleasure and pain subspecies). There may be other species, corresponding to so-called cognitive phenomenology, agentive phenomenology, and so on (see Kriegel 2015). This genus-species structure introduces another ambiguity in the question 'What is the value of consciousness?' To see why, suppose that fish and birds are valuable but amphibians, reptiles, and mammals are not. If someone asks whether animals are valuable, in one sense we should answer positively, since two species of the genus Animal are valuable; but in another sense we should answer negatively, since being an animal is not as such valuable. Thus when we ask about the value of a genus, we must distinguish two interpretations of the question: one about the existence of a valuable species and one about the value of the genus as such.

This is directly relevant to our purposes, because asking about the value of phenomenal consciousness could display the genus/species ambiguity just mentioned. Indeed, according to Andrew Lee (2018), algedonic phenomenal consciousness has intrinsic ethical value, but phenomenal consciousness does not as *such* have intrinsic value, as many other types of phenomenal consciousness are value-neutral. The point is that, because phenomenal consciousness is a genus, the question 'What is the value of phenomenal consciousness?' is also ambiguous as between a 'generic' and a 'specific' reading:

- What is the value that phenomenal consciousness *as such* has?
- What is (or are) the value (or values) that any *species* of phenomenal consciousness has (or have)?¹

Because phenomenal consciousness is a genus with several species, then, our eight questions about the value of phenomenal consciousness are really sixteen questions.

With the eight questions about the value of access consciousness (which is not in any obvious way a genus), they make up *twenty-four* different elaborations of 'What is the value of consciousness?'

3. The Having of Value by Consciousness

An assertion of the form 'Consciousness is V' (where V stands for a value property) has a subject-predicate structure. This structure may suggest that 'consciousness' denotes a substance (Descartes' *res cogitans*?) and 'V' a first-order property instantiated by that substance. But this cannot be right, since one need not believe in consciousness as a substance to find consciousness valuable. Many philosophers take the noun 'consciousness' to denote a property – the property of being conscious. These philosophers, too, should be able to make claims about the value of consciousness.

If the term 'consciousness' denotes a property, then the form 'Consciousness is V' can still be apposite, namely if 'V' denotes a *second-order* property that consciousness instantiates. But this too is problematic. At least for some substitution instances of V, properties will just not be the right kind of things to have V. Consider truth, an intrinsic epistemic value. It would be incoherent to say that this value is instantiated by the property of being conscious, given that properties are not truth-apt entities. Many putative entities are: beliefs, judgments, propositions, statements, assertions, and more. But a property is not the kind of thing that can be true or false. So this interpretation of 'Consciousness is V' is unsuitable as well.

How *should* we interpret such assertions? I propose that we see 'Consciousness is V' as claiming that anything that instantiates the property of being conscious will perforce also instantiate the property of being V – and will do so *because* it instantiates the property of being conscious. On this interpretation, the 'deep grammar' of 'Consciousness is V' is this:

- For any x, if x is conscious, then (i) x is V and (ii) x is V in virtue of x being conscious.

Suppose I claim that phenomenal consciousness has instrumental aesthetic value. Then what I am asserting is really this: Anything which is phenomenally conscious has, in virtue of that fact, instrumental aesthetic value. Accordingly, the question ‘What is the value of consciousness?’ really comes to the following: What are the value properties V_1, \dots, V_n , such that for any x , if x is conscious, then (i) x has V_1, \dots, V_n and (ii) x has V_1, \dots, V_n in virtue of being conscious?

There is still an ambiguity in this question, however. The question is effectively about grounding: about what holds in virtue of what. In the theory of ground, a distinction is sometimes drawn between *full* and *partial* grounds (Fine 2012). For when we say that p grounds q , this may mean either (a) that q holds *in part* because p holds or (b) that q holds *entirely* because p holds. In (a), a full specification of q ’s ground would have to cite some r that, when combined with p , ensures that q . In (b), it is p by itself that ensures that q . For instance, we may say that knowledge is grounded in belief, meaning that belief is a *partial* ground of knowledge. But we may also say that belief is grounded in tracking relation R and mean that R is the *full* ground of belief. Thus when we ask whether one thing is grounded in another, we must specify whether we are asking about full or partial grounding.

This distinction is relevant in our case, because an item’s being conscious may serve as a mere *component* of the grounds for that item’s value, or it may constitute those grounds *all by itself*. Accordingly, when we ask whether an item is valuable in virtue of being conscious, we must specify whether we mean ‘in part in virtue of being conscious’ or ‘entirely in virtue of being conscious.’

I closed §2 by noting twenty-four different questions folded into ‘What is the value of consciousness?’ If claims about the value of consciousness are really grounding claims, and grounding claims can be either full- or partial-grounding claims, then in each case we may ask either about consciousness’ role as full ground or about its role as partial ground. Thus the formulation ‘What is the value of consciousness’ hides at least *forty-eight* different questions – each of capital importance in its own right.

This factorization of our original question into forty-eight is obtained by applying a series of disambiguations: (a) intrinsic vs. instrumental value, (b) epistemic vs.

prudential vs. moral vs. aesthetic value, (c) phenomenal vs. access consciousness, (d) genus vs. species, and (e) full vs. partial ground. I leave it as an exercise to the reader to formulate explicitly all forty-eight questions. What I would like to suggest is that a comprehensive theory of the value of consciousness would consist in the conjunction of answers to all forty-eight questions!

Some value claims are more ambitious than others. A full-ground claim is more ambitious than a partial-ground claim, an intrinsic-value claim is more ambitious than an instrumental-value claim, and so on. But particularly ambitious are claims to the effect that not only consciousness has value *V*, but *only* consciousness has *V*. It is one thing to say that species *S* of phenomenal consciousness has value *V*; it is another to say that the only way for *V* to show up in our world is thanks to *S*. This kind of 'unique ground' claim is particularly ambitious. And when philosophers have claimed that there is no value in a world without consciousness (Seager 2001), what they have asserted is that for *any* value *V*, consciousness is the unique ground of *V*. Here we come close to the most ambitious thing that can be said about the value of consciousness.

PART II. SOME (POTENTIAL) ANSWERS

Moving forward I will ignore access consciousness and focus entirely on phenomenal consciousness. The next three sections are dedicated to airing some theses of potential interest about the epistemic, ethical, and aesthetic values of (various species of) phenomenal consciousness.

4. The Epistemic Value of Consciousness

The epistemic value of consciousness has been the most explored among philosophers.² In particular, it has sometimes been argued that some perceptual beliefs are justified by perceptual experiences and are so in virtue of those experiences' phenomenology (Pryor 2000 fn37, Smithies 2014, Brogaard and Chudnoff forthcoming).

This is not the place to consider the full case for such a claim, but for a taste of the type of consideration relevant, consider the following scenario.

(Dark Room) You wake up from a groggy, unexpected nap and find yourself in a pitch-dark room that feels unfamiliar. A warm voice greets you and then asks you whether you (a) believe that there is a chair in the room, (b) disbelieve that there is a chair in the room, or (c) suspend judgment about whether there is a chair in the room. After you answer, the lights come on, and you vividly experience a perceptual phenomenology as of chair in front. The warm voice comes on again and asks whether you (a) believe that there is a chair in the room, (b) disbelieve that there is a chair in the room, or (c) suspend judgment about whether there is a chair in the room.

Intuitively, the doxastic attitude it is most rational (or most prima facie justified) for you to take the first time you are asked about the chair is to suspend judgment; but the second time you are asked, with the lights on, it is more rational for you to believe. The rationality in question is clearly epistemic, since what you are asked to do is make up your mind about *what is true* (or *what is the case*).

What explains the difference in which doxastic attitude is rational for you before and after the lights come on? One natural explanation is that after the lights come on you enjoy a phenomenology as of perceiving a chair that you did not before – and it is the occurrence of this phenomenology that rationalizes your chair belief. (Objection: if you imagine a chair, you have the same chair-y phenomenology, but that would not rationalize the belief that there is a chair in the room. Response: a phenomenology as of perceiving a chair has more to it than just a chair-y phenomenology; it also involves a perception-y phenomenology, that is, a phenomenology of commitment to the reality of what one is aware of. The latter phenomenal element – sometimes called ‘phenomenal force’ in the relevant literature – is a sort attitudinal phenomenology that distinguishes perceptual from imaginative experience.) If this explanation of the difference in which doxastic attitude is rational is better than all alternative explanations, that would provide substantial evidence for an epistemic value proper to perceptual consciousness.

One alternative explanation might be that your chair belief is justified by the reliability of your perceptual mechanisms, at least as far as visual chair-detection in well-

lit rooms is concerned. And reliability has nothing to do with consciousness. The problem with this, however, is that the intuition about the change in which attitude is rational does not disappear when we replace you with a brain in a vat. Certainly an envatted phenomenal duplicate of you who came to believe that there was a chair in the room *after* the lights came on would be epistemically blameless in a way an envatted duplicate who believed in the chair *before* the lights came on would not be. Yet the duplicate's chair-detection mechanism is entirely unreliable (leading as it does exclusively to false beliefs), and is equally so when the duplicate has a phenomenology as of perceiving a chair and when it has a phenomenology as of pitch dark.

Another alternative explanation might appeal not to the reliability of your chair-detection mechanisms, but to your (justified) *belief* in the reliability of these mechanisms – a (justified) belief that your envatted duplicate shares. The problem here is that the reliability belief is present even before the lights come on; so clearly, it is not *sufficient* to rationalize a chair belief. The reliability belief seems to be 'epistemically charged' (like a battery is charged) throughout the scenario, but to become epistemically *operative* (in some sense) only when plugged into the phenomenology of perceiving a chair. Thus this 'alternative' explanation seems to invoke perceptual phenomenology too; it just avails itself of an *additional* element in the explanation, namely, the reliability belief. This addition may diminish, but does not vacate, the epistemic value of perceptual phenomenology. More precisely, it appears to 'downgrade' perceptual phenomenology from full to partial ground of epistemic value (but not to *remove* perceptual phenomenology from the ground).

There are probably other alternative explanations that a thorough argument for the epistemic value of perceptual phenomenology would have to address. Still, the above discussion provides a sense of the kinds of consideration that support the attribution of epistemic value to consciousness.

The epistemic-value claim made here is already controversial, with some epistemologists wondering why the fact that one experiences a certain phenomenology and not another should generate justification for anything (see, e.g., Ghijzen 2014). But even if we grant it, it is noteworthy that the claim is limited in at least two ways. First, it concerns the epistemic value not of phenomenal consciousness as such, but merely of one species of it – the perceptual species. Secondly, by the lights of most views, it

concerns only *instrumental* epistemic value. Certainly veritists will claim that epistemic rationality is valuable only insofar as it is instrumental in promoting true belief; but even such proposed alternative intrinsic values as understanding go beyond epistemic rationality.³ Now, the fact that the relevant value is merely instrumental, and attaches only to perceptual consciousness, does not by itself mean that it is inconsequential. Pautz (2017), for instance, argues that this alone shows that most versions of physicalism are false (roughly: because they cannot explain how physical properties that differ in the physically tiniest ways can be differ so dramatically in terms of value). Still, it would be interesting to consider whether more ambitious epistemic-value claims could be made on behalf of consciousness.

As far as generalizing from the perceptual case, one general strategy is to push broadly perceptual models of certain mental phenomena and then use them to extend perceptual phenomenology's epistemic role to other species of phenomenal consciousness. Thus, philosophers who hold that emotional experiences are value perceptions may claim that some value beliefs are *prima facie* justified by emotional phenomenology in the same way some empirical beliefs are *prima facie* justified by perceptual phenomenology (cf. Elgin 2008); philosophers who endorse a perceptual model of intuition, and hold that intuition exhibits a kind of cognitive phenomenology, may claim that some (e.g.) mathematical beliefs are *prima facie* justified by the phenomenology of intuition in the way empirical beliefs are *prima facie* justified by perceptual phenomenology (Chudnoff 2011); and so on. Similarly, it is possible to hold that the phenomenology of episodic memory provides *prima facie* justification for beliefs about the past, the phenomenology of imagination *prima facie* justification for modal beliefs, and so on. If every type of phenomenal consciousness could be shown to exhibit some kind of epistemic value, and if some such kind could be shown to be exhibited by *all* species of phenomenology, that value could be seen as belonging properly to phenomenal consciousness *as such*.

Even then, of course, this would still be merely instrumental value. What are prospects for finding some *intrinsic* epistemic value in consciousness? One option here is to argue that epistemic rationality or *prima facie* justification is epistemically valuable in and of itself. But this is highly dubious, and anyway it would be far better, from a dialectical standpoint, to argue that consciousness is implicated in more independently

motivated and widely recognized intrinsic epistemic values, notably truth and understanding.



Obviously, we have no doxastic or otherwise truth-apt state in our psychological repertoire such that its being conscious guarantees its being true. To that extent, consciousness' prospects for veritistic value may seem unpromising. On some views of consciousness, however, every conscious experience involves a built-in awareness of its occurrence. If one held that the relevant awareness is for some reason guaranteed to be veridical (Kriegel 2009 Ch.4, Kidd 2011), this would cast conscious experience as veritistically valuable after all: the occurrence of every conscious experience brings with it the occurrence of a veridical representation.

In addition, it is noteworthy that although the broadly Cartesian idea that all introspection is infallible is roundly rejected in contemporary philosophy, several philosophers have argued that there exists a restricted type of introspection which is in principle immune to error (Gertler 2001, Chalmers 2003, Horgan and Kriegel 2007, Giustina forthcoming). If this is right, and if phenomenally conscious states are the only things at which the relevant type of introspection can be directed (as all these authors will accept), then a mental state's being conscious is a precondition for its being infallibly known. This is something a veritist should appreciate.

At the same time, it is not clear that this would cast consciousness as really *intrinsically* epistemically valuable. Suppose God made sure that we could not misperceive or misjudge about zebras, though elsewhere we be thoroughly fallible. That would not quite give *zebras* intrinsic epistemic value. It would rather be zebra perceptions and judgments that would be intrinsically valuable. Zebras would have *instrumental* value insofar as they *enabled* infallible perceptions and beliefs, but it would be the relevant perceptions and beliefs that would be intrinsically epistemically valuable.⁴

A more promising approach, I suspect, might be found outside the veritistic fold – in particular in the idea that there is a kind of intrinsically valuable form of *understanding* that is fully or partially grounded in conscious experience. An argument to that effect is provided by Bourget (2017).⁵ The argument proceeds in two steps. The

first notes a common theme in theories of understanding within epistemology and philosophy of science, namely, that understanding requires *grasping*. These views differ on just what grasping is, and what exactly needs to be grasped for something to be understood; but they agree that without grasp there is no understanding (Gordon 2017). Bourget gives the following illuminating example: when told by trustworthy authorities that the sun is approximately 1,300,000 bigger than the earth, we believe this (correctly) but in a sense have no grasp of what this really means; when later we are told that that is like the difference between a basketball and an apple seed, something happens with us which can rightly be described as an epistemic achievement. But the achievement is a peculiar one: we do not have any further *evidence* for the proposition that the sun is approximately 1,300,000 bigger than the earth when told about the seed and the basketball. Rather, we gain a measure of insight into what it really means that the sun is approximately 1,300,000 bigger than the earth – we grasp the fact in a way we did not previously. The second step of Bourget’s argument is to ask what distinguishes believing that *p* without grasping *p* and with grasping *p*, and argue that the difference has to do with phenomenal consciousness: we believe *with* grasping when our belief is accompanied by a phenomenal experience with the content *p*; we believe *without* grasping when our belief is not so accompanied. Bourget argues that this explanation is better in various respects than functionalist and other explanations.

One limitation of Bourget’s discussion is that he does not tell us what kind of experience makes for grasping, and indeed seems to suggest that *any* experience would – something which is rather implausible. (If I believe that the weather is nice while phenomenally experiencing a hope that the weather is nice, that does not seem to make for grasping of the weather’s niceness.) But we may find help on this point in Chudnoff’s (2012) notion of *presentational phenomenology*: a kind of phenomenal characteristic whereby a mental state not only represents *p* but also makes it seem to the subject that she is aware of the *truthmaker* of the proposition *p*. Importantly, for Chudnoff presentational phenomenology can occur in a variety of conscious experiences, including perception, episodic memory, introspection, and intuition. The difference is just in the kind of awareness of the putative truthmaker involved: while a perceptual experience of *p* may make one seem to be sensorily aware of the truthmaker, an intuition that *p* may make one seem to be intellectually aware of the truthmaker. Using this notion of presentational phenomenology, we may surmise that

the difference between believing that p with grasping and without grasping is that although both beliefs represent p , only the former involves a presentational phenomenology whereby the subject seems to be aware of the truthmaker of p .

Bourget and Chudnoff are concerned with propositional grasping, but in certain areas much emphasis has been put on grasping the *natures* of certain things (often: of properties). To appreciate the idea, consider the curious case of Unlucky Luke:

(Unlucky Luke) Luke is a guy like you and me, but his life thus far has suffered from a double misfortune. First, whenever Luke has come across red objects, it was through a thin mist that made them appear pink to him. Second, whenever Luke has come across a pink object, there was a strong red light shining on that object, unbeknownst to Luke, that made it look red rather than pink to Luke.

Unlucky Luke has a massively wrong conception of the extensions of red and pink: he thinks fire engines are pink and ballet shoes are red. At the same time, there is a sense in which Luke's grip on *what red is* and *what pink is* is as good as yours and mine. He has the same grasp of the *nature* of red and pink as you and I. This is the kind of grasp that Frank Jackson's (1982) Mary, for instance, lacks. Despite her flawless conception of the extensions of red and pink, the intuition is that Mary has no grasp of the natures of red and pink (Conee 1994).

Arguably, this grasp of natures requires conscious experience as well. If we suppose that upon release it turns out that Mary is blindsighted for red, such that her visual system can unconsciously process information about red objects, but she can never enjoy a conscious perceptual experience as of red, then the intuition persists that she still does not know the nature of red.

I belabor these points because they sketch a potential non-*veritistic* defense of consciousness' intrinsic epistemic value. If (a) understanding is an intrinsic epistemic value, (b) grasping is an ineluctable component of at least certain types of understanding, and (c) presentational phenomenology or some such phenomenal feature is necessarily involved in grasping, then (d) there is a species of phenomenology that partially grounds intrinsic epistemic value.

5. The Ethical Value of Consciousness

In §1, I distinguished between prudential and moral value, as two kinds of ethical value. This particular bookkeeping regime is not what matters here. What matters is that when we describe a life as a good life, we may have in mind a life good for the one who lives it, or a life that is morally good. Colloquially: either the life of a *happy* person or that of a *good* person. The relationship between these two notions may be intimate (Annas 1993), but they are clearly conceptually distinct. The contemporary literature on well-being could be seen as addressing the question of what makes life good for the one who lives it, while the vaster literature on moral philosophy, featuring notably consequentialist, deontological, and virtue-ethical theories, addresses the question of what makes a life morally good. Our question is: What is the role of consciousness in each?

5.1. *The prudential value of consciousness: Consciousness and well-being*

A common if unloved classification divides theories of well-being into three general approaches: hedonism, desire-satisfaction, and objective-list theories. In what follows, I will consider the role consciousness might play in grounding prudential value within each approach.

According to hedonism, a life's goodness is determined by the distribution in it of hedonically valenced experiences, notably pleasure and pain: the more pleasure and the less pain, the better the life for the one who lives it. This delivers fairly straightforwardly an intrinsic prudential value for consciousness. All we must add is that pleasure and pain contribute to well-being *in virtue of their phenomenal character* – what it is like to experience them (Bramble 2016). On its face, this extra claim is rather plausible. Imagine a functional duplicate of yours with the following peculiarity: its mental states that play the pleasure role always lack phenomenology and are entirely unconscious, while its mental states that play the pain role are phenomenally experienced just as yours are. From a hedonic standpoint, this ill-fated duplicate seems to lead a rather sorrowful existence, and certainly not to live a life as good as yours. (Note that we can appreciate this point while staying silent on the question of whether the duplicate's functionally equivalent states count as pleasures and pains.)

It might be objected that what makes our pleasure of value to us is not simply that pleasure *has* a certain phenomenology, but that we *like* this phenomenology and *want* to experience it. Indeed, Heathwood (2007) has argued that pleasure is in fact *nothing but* a compound mental state involving a desire directed at a phenomenology.⁶ Now, this does not *by itself* undermine the role of consciousness in well-being, insofar as one might still require that the relevant desire be conscious. Still, I find it strange to treat our desire for pleasure phenomenology as explanatorily prior here. After all, surely we like and want the phenomenology of pleasure because of *what that phenomenology is like*. It would be odd to treat as brute and inexplicable the facts that we like and desire pleasure phenomenology and dislike and are averse to pain phenomenology. These facts do not seem like rock bottom at all. They seem *obviously* explicable, namely by the nature of pleasure and pain phenomenology. It is because pleasure feels *like this* that we want it. If anything, then, it is the facts about how pleasure and pain feel like this that ground hedonic well-being.

If all this is right, then hedonism leads to the thesis that intrinsic prudential value is fully grounded in a species of phenomenal consciousness. The relevant species may be construed as narrowly algedonic (pleasure and pain), or it may be widened to include a range of affectively valenced experiences, such as hope, joy, awe, satisfaction, nostalgia, gratitude, affection, contentment, as well as a variety of basically ineffable shades of 'good mood.'⁷ Some of these experiences it may be awkward to describe as 'pleasures,' though all seem aptly described as 'pleasant' or 'pleasurable.' Why the noun is awkward but the adjectives are felicitous is something we need not spend time on here. To bracket terminological issues, we may call 'affectivism' the view that well-being is fixed by the distribution of affectively valenced phenomenology in it, and leave it undecided whether affectivism is a *version* or an *extension* of hedonism. What matters to us here is that affectivism grounds intrinsic prudential value in a more generic kind of phenomenal character than narrowly algedonic phenomenology.⁸

Now, hedonism (including affectivism) faces two extraordinary objections, which have turned many ethicists against it. These are the 'philosophy of swine' and 'experience machine' objections. The former may be put as follows: hedonism cannot vindicate the strong intuition that the pleasures of reading poetry, having insights, and so on are somehow *more valuable* than the pleasures of orgasms and strawberries, despite the latter being phenomenologically more intense than the former. The

'experience machine' objection is that hedonism cannot vindicate the intuition that it would be rational for us to decline an invitation to enter a machine that would simulate exactly the same life we would otherwise have but add to it a bit of extra pleasure (one more strawberry!).

My own view is that the hedonist actually *can* vindicate the swine intuition (Crisp 2006, Bramble 2016),⁹ and that she has some resources to debunk the experience-machine intuition (Silverstein 2000, Crisp 2006, De Brigard 2010). It is worth noting, in any case, that that even if we accept the experience-machine objection at face value, all it really shows is that experience is not *sufficient* for well-being, or that it does not determine well-being *all by itself*. The objection does nothing to undermine the *necessity* of experience for well-being, the idea that experience makes a crucial *contribution* to well-being. For suppose you had to choose between the following two options: either (a) you enter the experience machine and lead a life phenomenally approximately indistinguishable from the one you would be leading outside the machine; or (b) you stay outside the machine but lose all phenomenal consciousness, carrying on your 'public' life entirely undisturbed – but as a complete zombie. It would seem folly to choose (b) in this scenario. Thus the fact that in the original experience-machine scenario you are guaranteed a continued experiential life outside the machine is crucial in sustaining the intuition that it is preferable for you not to plug in. At most, then, the experience-machine argument shows that (affectively valenced) experience is only a *partial ground* of intrinsic prudential value.

This is independently plausible. For consider the following choice: either (a) you submit today to a medical procedure that would irreversibly zombify you overnight (though it will leave your cognitive and motivational architecture as functionally unchanged as possible) or (b) your life from now on will skew consistently towards the unpleasant (it is not that you will be forever tortured, but just that you will experience a lot of frustration and pain and only very rare, and very mild, joy). I think most of us will have no difficulty choosing (b). At least as long as the pain is not too overwhelming, most of us find it better to live with it than to end our conscious life. This seems to suggest that we value being conscious as such, regardless of the hedonic quality of our consciousness.¹⁰ If this is right, then prudential value is grounded partly in algedonic phenomenology and partly in phenomenal consciousness as such.



Partly under the pressure of the experience-machine objection, many ethicists today prefer a desire-satisfaction theory, according to which a person's well-being is a function of her desires being satisfied. It is not the distribution of pleasure vs. pain that matters, but the distribution of satisfied vs. frustrated desires. Importantly, the notion of satisfaction at play here is not subjective but objective: a person's desire being satisfied is not a matter of that person experiencing a phenomenal feeling of satisfaction, but of the desire's content-fixing satisfaction conditions being met (ditto for 'frustration'). The person's desire can be satisfied in this objective sense even if the person is altogether unaware that it is, indeed if the person is long dead (e.g., my desire that the Israeli-Palestinian conflict be peacefully resolved one day may well be satisfied after my death). It is this notion of satisfaction that helps the desire-satisfaction theorist avoid the experience-machine problem: in the experience machine, we have the *impression* that our desires are being satisfied, and accordingly experience feelings of satisfaction, but our desires are not *really* satisfied (e.g., I may want to go back tonight to the amazing restaurant from yesterday, when in reality I have been tube-fed indeterminate mash ever since plugging into the machine 20 years ago). What matters for one's well-being, on this view, is that one's desires be genuinely satisfied; the presence or absence of feelings of satisfaction (and their positive affective phenomenology) is immaterial.

In addition to being supported by an array of arguments (Heathwood 2016: 137-8), taking satisfied desires to be the only things of intrinsic prudential value has a pleasing symmetry with taking true beliefs to be the only things of intrinsic epistemic value. And like the idea that true beliefs are the only things of intrinsic epistemic value, it does not seem to assign special value to phenomenal consciousness (Levy 2014). Even if we suppose, rather implausibly, that all desires are phenomenally conscious, and moreover that their satisfaction conditions are fully fixed by their phenomenal character, what matters for prudential value on this theory is not just the *having* of desires, but their being satisfied. Prudential value comes into the picture when we move from desires in general to satisfied desires. And what makes a desire satisfied has nothing to do with consciousness, but with the objective facts about the world (except of course when consciousness is part of what the desire is a desire *for*).

My own view is that desire-satisfaction theory is too implausible, though, to present an important challenge to the role of consciousness in well-being.¹¹ This has been demonstrated to my satisfaction by Charles Siewert (1998 Ch.8, forthcoming Ch.16), and here is the heart of the demonstration. Suppose you are offered the following deal by an omnipotent deity: if you submit to the zombification procedure that retains your cognitive and motivational architecture, the deity will see to it that an overwhelming majority of your current desires be satisfied. Should you accept the deal if you want to make your life better (for you)? Intuitively, you should not. Intuitively, your life will not be better for you if you submit to this procedure; it will be much worse. In fact, there is a perfectly good sense in which *your* life will be no more (Cutter 2017). For suppose you are offered a variant of the deal where after a week in a state of zombiehood, you would be killed and replaced with another zombie cooked up in a Petri dish (and functionally indistinguishable from your bezombied self, naturally), but your current desires will continue being methodically satisfied one by one. It is entirely unclear what reason you have to find this new variant of the deal any less appealing than the previous one. Both offers, I claim, are really terrible *for you*: far from making your life better for you, they would end your life in any prudentially significant sense of 'life.' (It is important here to screen off intuitions about whether it would be *morally* good to accept the deal. You might affect a great deal of good in the world if you thus sacrifice yourself to secure the prospective satisfaction of all your geopolitical and macroeconomic wishes. But that would still be a *sacrifice*: your own well-being would be the price you would heroically pay for things going well for others.)

The third type of theory of well-being (in common taxonomies) is the objective-list theory. The general idea is that a life is good for the one who lives it to the extent that it features as many items from a privileged set of goodness-conferring elements. Strictly speaking, however, hedonism and desire-satisfaction theory also satisfy this general characterization. What marks off objective-list theories is two main features, one captured in 'objective' and one in 'list.' The 'objective' part is that the goodness-conferring elements must be such independently of our subjective reactions to them: it is not because we enjoy or want them that they make our lives better. The 'list' part intimates a strong tendency toward pluralism: we expect the goodness-conferring elements to be several, rather than to reduce to a single source. Thus a typical objective-list theory would cite love and friendship, achievement and self-realization,

play and free time, appreciation of beauty, and a number of other elements the very existence of which in one's life – regardless of one's attitude toward them – makes it a better life.

Objective-list theories may seem initially to be least accommodating for a special role of consciousness in well-being. But in fact most theories in this vein do cite affectively valenced experiences in their list of objective goodness-conferring elements; it is just that (i) it is the having of those experiences that they consider good, rather than our attitudes toward them, and (ii) their lists include also other, non-experiential elements. Thus, Fletcher (2013) includes happiness and pleasure in his objective list, while Murphy (2001) includes happiness, 'inner peace' (which is presumably feels quite nice!), and aesthetic experience (presumably of the pleasant variety to do with beauty rather than the unpleasant one to do with ugliness). To be sure, the fact that for such objective-list theories valenced experiences are only one of many elements responsible for well-being diminishes considerably the prudential value of consciousness; but it does not annihilate it. More precisely, it casts affectively valenced experience as a partial rather than full ground of well-being.

Furthermore, we should keep in mind the coherence of a sort of limit case of objective-list theory that includes on its list *only* conscious experiences, though not only algedonic and affective ones. The items on the list are subjective, but it is an objective fact that they belong on the list – it is not because we want or enjoy these experiences that they belong there, but because of what these experiences are like in themselves. In elaborating this kind of 'list-experientialism,' van der Deijl (forthcoming) cites self-understanding and novelty as experiential goods. Other philosophers have belabored the value of *cognitive* phenomenology. Galen Strawson (2011: 299), for instance, argues that without cognitive phenomenology life would be 'pretty boring' (see also Chudnoff 2015 Ch.3). And as noted above, we may well value having phenomenology as such (so long it is not too agonizing). Thus one may draw a longish list of valuable experiences that go beyond the algedonic and affectively valenced, and ground well-being in it. In this limit case of objective-list theory, experience regains its status as *full* ground of intrinsic prudential value.¹²

I conclude that unless one adopts either (a) a desire-satisfaction theory or (b) a specially idiosyncratic objective-list theory purged of all experiential elements,

consciousness likely plays a crucial role in well-being, and may even turn out to be the full ground of intrinsic prudential value, namely if either hedonism or an experientialist version of objective-list theory bear out. My own sympathies lie with a version of list-experientialism that features only two items on its list: affectively valenced experience and phenomenal experience as such. To that extent, I am inclined to think that prudential value is fully grounded in phenomenal consciousness.

5.2. The moral value of consciousness: Consciousness and morality

As noted, modern moral philosophy has been dominated by three main approaches: consequentialism, deontology, and virtue ethics. Here I will treat these as approaches to what makes a life morally good.¹³ Let us consider, then, the outlook within each framework for a role for consciousness in moral value.

For present purposes, the best way to characterize consequentialism is as the view that a life is morally good to the extent that it maximizes well-being in the world – not just one’s own, but others’ as well. When one combines this with hedonism about well-being, one obtains Bentham/Mill-style classical utilitarianism: the morally good life is one that maximizes pleasure and minimizes pain. But consequentialism can also be combined with a desire-satisfaction account of well-being, the result being Hare’s (1981) ‘preference utilitarianism,’ or with objective-list accounts, yielding any number of pluralist consequentialisms (see already Brentano 1889 and Moore 1903).

Because consequentialism effectively embeds theories of well-being in its conception of ‘the consequence,’ our conclusions about prudential value at the end of §5.1 apply *mutatis mutandis* to consequentialism about moral value: classical utilitarianism and most versions of pluralist consequentialism (including Brentano’s and Moore’s, both of whom cite pleasure among the intrinsic goods we should promote) will assign consciousness (in particular algedonic phenomenology) a central place in grounding moral value; preference utilitarianism and experience-purged versions of pluralist consequentialism will not.

Deontological approaches to morality emphasize motivation over consequences: they recommend acting from the motive of duty rather than acting for the promotion of well-being. There is any number of views one could have about what our moral duty is, but the paradigmatic view here is surely Kant’s: our duty is to always treat human

beings as ends in themselves rather than as mere means to our own ends (or for that matter someone else's ends). There is something about the presence of another person in a situation that, for Kant, puts the breaks on how we would otherwise naturally act – the person 'exacts respect for himself from all other rational beings in the world' (Kant 1797: 186-7). That in virtue of which a person exacts respect and merits being treated as an end is what Kant calls 'dignity.' Our fundamental duty, then, is to respect and treat as ends all creatures possessed of dignity; all our other, more 'local' duties derive essentially from this one.

What role might consciousness play within the deontological framework? I think the key question here is: What gives some things dignity and not others? It seems odd to think that dignity is a brute, inexplicable attribute that attaches itself to some entities and not others in a completely arbitrary way. More plausibly, things have certain empirical properties that ground their dignity – be these genetic, psychological, or other empirical (i.e., non-evaluative) properties. We have duties to other people, but not to rocks – and this must have to do with the factual differences between people and rocks. Elsewhere, I have argued that the ground of dignity is precisely phenomenal consciousness (Kriegel 2017). On this view, an entity exacts respects and merits treatment as an end just if it is a phenomenally conscious creature. (It follows that we have duties toward not only human beings but all conscious beings, including conscious nonhuman animals.) If consciousness is the ground of dignity, and dignity the ground of moral value, then by transitivity of ground, consciousness is the ground of moral value.

It might be objected that this fits poorly with Kant's idea that dignity attaches to rational beings capable of setting ends. But the Kantian idea faces a dilemma: could an unconscious automaton be a rational end-setting being? If not, then (a) it would seem that rational end-setting status implies consciousness. But if a rational end-setting automaton is possible, then (b) intuitively, there is no particular reason for us to treat it as an end rather than a means. We can all envisage a future in which household robots are treated as means for the performance of household chores. Even if such robots are developed to the point that they could set their own ends, as long as they remained certified zombies, we would feel no more pressure to treat them as ends than we do our current-day Roombas. Both those future robots and today's Roombas are characterized by one simple fact: *there is nobody home*. For these reasons, I am

convinced that Kant himself had horn (a) in mind, and so would accept consciousness as at least a partial ground of dignity ('partial' because end-setting capacity would be required as well).

As an aside, let me state that one could always adopt a hybrid ethical theory combining deontological and consequentialist elements. Indeed, my own ecumenical moral maxim would be: Always maximize well-being in the world consistently with treating all creatures possessed of dignity as ends rather than mere means. Given my views on well-being and dignity, this comes to the following:

(Ecumenical.) Always maximize positively valenced (and minimize negatively valenced) experiences in the world consistently with treating all phenomenally conscious creatures as ends rather than mere means.

This maxim is doubly consciousness-invoking: (i) insofar as it identifies the promotion of positively valenced phenomenology as the default goal of moral action and (ii) insofar as it identifies the possession of phenomenal consciousness as the only factor constraining the pursuit of that default goal.

As noted, I am tempted by a moral theory organized around something like *(Ecumenical)*. According to this, there are two independent sources of moral value in our world, well-being and dignity, and both are empirically grounded in phenomenal consciousness: well-being in valenced phenomenology, dignity in phenomenal consciousness as such. Thus on the moral theory I am most tempted by, consciousness lies at the grounds of moral value.



The last major approach to moral value is virtue ethics. Here what makes a person's life morally good is his or her actively maintaining and reliably exercising stable dispositions of a certain sort – the virtues. That is what we call a *good person*: a person on whose generosity, modesty, kindness, empathy, trustworthiness, warmth, fairness, and courage can be safely relied upon. The morally good life is the life of such a person.

What role does consciousness have in this framework? Well, whatever role it has in the virtues. Now, being dispositions, the virtues are standing conditions of personality, not occurrent events as conscious experiences are. Still, consciousness may

well show up in the identity conditions of these dispositions. In general, dispositions individuate by their *triggering* and *manifestation* conditions. To take a toy example, a vase's disposition to break when dropped (roughly: 'fragility') has dropping for triggering condition and breaking for manifestation condition. Generosity, too, has characteristic triggering and manifestation conditions. The triggering condition of generosity is typically the presence of someone in need. Arguably, now, a zombie duplicate of a person in need either (a) does not have needs the meeting of which has any moral significance or (b) does not have needs at all, in the proper sense of 'needs.' So consciousness appears to show up in the triggering conditions of generosity. As for the manifestation conditions, the central such conditions are surely behavioral and involve the ceding of resources (money, time, energy) to the person in need. But in the Aristotelian tradition, the generous person is she who not only behaves in the right way but also thinks and feels the right way. Indeed, in the *Nicomachean Ethics* Aristotle himself puts feeling at the heart of virtue: 'to feel [what we feel] at the right times, with reference to the right objects, towards the right people, with the right motive, and in the right way, is ... characteristic of virtue' (6.1106b). Thus generosity is individuated partly in virtue of the kind of magnanimous or giving feeling it manifests in. There is a subtle phenomenology of magnanimity that one must be able to appreciate for one to understand what a manifestation of generosity is – and therefore what generosity itself is.

Obviously, I do not offer the foregoing as a robust argument that consciousness shows up both in the triggering and manifestation conditions of generosity. The considerations I have aired are far too superficial and preliminary. All the same, they serve to *illustrate* how consciousness might be claimed to show up in the identity conditions of a virtue – to be part of what makes it the virtue it is. Now, generosity may be unusual in having consciousness in its triggering conditions. Courage, for instance, is triggered by conditions of danger, and these can occur in the complete absence of consciousness. Still, even courage plausibly involves, as part of its *manifestation* conditions, a certain phenomenology of resolve in the face of danger. And from an Aristotelian perspective, we may expect some experiential component in the manifestation conditions of all or most virtues.¹⁴

For a virtue ethicist, then, moral value is grounded in the virtues, and certain species of consciousness (a phenomenology of magnanimity, a phenomenology of

resolve, and so on) are part of what individuates the virtues. It follows that certain species of consciousness are part of what individuates the grounds of moral value. To that extent, they may be thought of as partial grounds of moral value. At the same time, it is surely a *small* part. Behavioral manifestations will always be more crucial to the identity conditions of the virtues than experiential manifestations, and many virtues' triggering conditions will not include consciousness at all.

In conclusion, we see that many mainstream views on both types of ethical value – prudential and moral – ground this value in phenomenal consciousness. Typically the claim is of merely partial grounding, and of grounding by only some species of phenomenal consciousness. But it is also typically a claim of *intrinsic* value (unlike the epistemic case, where the standard claim concerns merely instrumental value), and on the margins we have also encountered claims of *full* grounding (in hedonist and experience-list theories of prudential value and corresponding versions of consequentialism about moral value). Thus the ethical value of consciousness is by many lights quite important.

I mentioned at the opening of this essay that the value of consciousness has been a lively area of debate in medical ethics. A big reason for this is the dramatic emergence of very strong evidence of consciousness in patients in permanent vegetative state (PVS) relatively recently. The question before the ethicists is whether such patients should be granted 'full moral status,' which includes the right to life. The shared assumption seems to be that they should, provided the evidence suggests that they can lead lives worth living (i.e., lives that clear a certain threshold of goodness to those who live them). But this is where opinion split. Kahane and Savulescu (2009) argue that PVS patients should not be kept alive, on the grounds that their inability to communicate with others likely results in unbearable solitude. Here the background framework seems to be broadly hedonist: it is the pain caused by solitude, and its (speculated) overriding enormity, that grounds Kahane and Savulescu's verdict. This hedonist background seems to be shared by Peterson and Bayne (2018), who only call into question Kahane and Savulescu's empirical assumptions about the possibility of meaningful interaction in PVS. In contrast, Levy and Savulescu (2009) argue that phenomenal consciousness is not what matters to a patient's full moral status, but rather the patient's hope that many of her interests will be satisfied. This seems to presuppose rather a desire-satisfaction theory of well-being, something Levy (2014) is

explicit about. The dispute goes on, and mixes empirical questions about the nature of PVS patients' actual conscious lives (something we know relatively little about) and ethical questions about what makes a life worth living. To my knowledge, the dispute does not feature any disputants who rely on an objective-list theory.¹⁵

6. The Aesthetic Value of Consciousness

We have many experiences of beauty in life – aesthetic experiences occasioned by encounter with art, with nature, and even with abstract structures. But do we have experiences that are themselves beautiful? Do we have aesthetic experiences directed at *other experiences*? For my part, I find it hard to think clearly about what that would amount to. But when I widen my view from beauty to the great range of aesthetic values commonly cited in the philosophy of art, I find it easier to conceive of aesthetic values borne by experiences. A fine shade of moral emotion may seem *exquisite*, an insight that allows one to see a longstanding problem in a fresh and fecund light may seem *compelling*, and so on. At the same time, clearly many things other than conscious experiences can be exquisite or compelling, and to that extent the relevant experiences appear inessential to the *existence* of those aesthetic values. And it is also hard to know what exactly accounts for the exquisite or compelling aspect of these experiences. Certain experiences may, by their rarity and intricacy, exhibit aesthetic value; but it may well be their rarity or intricacy that grounds this value, rather than their being experiences.

Among the traditional aesthetic values, perhaps *the sublime* is the one most closely tied to consciousness.¹⁶ On the traditional conception of the sublime, the sublime is not simply whatever is especially beautiful. Rather, it is what is so incomprehensibly grand and imposing as to somehow elude our grasp and inspire in us something like awe. For my part, as I contemplate particularly vividly the explanatory gap and the apparent categorical difference between consciousness and the rest of the natural order, a certain theoretical type of awe descends on me, whereby consciousness seems sublime. (Note well: I experience this awe despite my credence in physicalism being higher than .5!)

The real potential for a central role of consciousness in aesthetic value seems to lie elsewhere: in the conditions for the very possibility of such value. What I have in mind is a kind of subjectivist or 'secondary-quality' account of beauty and other aesthetic values, an account that grounds a thing's aesthetic value in corresponding aesthetic experiences. On a traditional secondary-quality of color, say, an object being verdigris is a matter of it eliciting, or being disposed to elicit, verdigris-ish visual experiences in the right kind subject under the right kind of circumstance. By the same token, it is not unnatural to suggest that an object being beautiful is a matter of it being disposed to elicit a positive aesthetic experience in the right kind subject under the right kind of circumstance. This schema can be generalized to all intrinsic aesthetic values (in case there are others than beauty), and made more precise, as follows:

- (*) For any item x and intrinsic aesthetic value V_i , such that x is V_i , there is a kind of subject S , a kind of circumstance C , and an aesthetic experience \mathcal{A}_i , such that (i) x is disposed to elicit \mathcal{A}_i in S under C and (ii) x is V in virtue of the fact that (i).

This grounds aesthetic value in aesthetic experience, a species of phenomenal consciousness, though it does so in an importantly different way than the value-ground claims considered in this essay so far. (More on this below.) Presumably, the phenomenology of aesthetic experience is crucial here: if an object elicited in us internal states with the functional role of aesthetic experience, but without the distinctive phenomenology of aesthetic delight, most subjectivists would not be inclined to consider it beautiful.

One immediate problem with this approach to aesthetic value is that it does not seem to recover the evaluative or 'normative' aspect: what experiences an object is disposed to elicit is a purely factual, 'descriptive' matter. To this difficulty, two approaches can be taken.

One approach is to unpack 'right subjects' and/or 'right circumstances' in overtly evaluative terms. Thus, one might construe these as *ideal* subjects and/or circumstances. The result is a sort of 'ideal-observer theory': an object being beautiful is a matter of it eliciting a positive aesthetic experience in *ideal* subjects under *ideal* circumstances. Of course, what makes a subject ideal cannot be that she has *veridical* experiences of aesthetic value (and what makes a circumstance ideal cannot be that it favors the formation of such veridical representation), on pain of circularity. Some story

must be told about what makes subjects (and circumstances) ideal that does not invoke objects' aesthetic value. Here Hume's version of subjectivism points the way, with its account of five central elements that distinguish 'true critics' from ordinary subjects: 'Strong sense, united to delicate sentiment, improved by practice, perfected by comparison, and cleared of all prejudice' (Hume 1757: 229).

The other option for recovering the evaluative dimension of beauty is to revise the account so it grounds beauty not in the experiential responses subjects *actually* have but those it would be *right*, or *fitting*, for them to have. This is a sort of 'fitting-attitude account' of aesthetic value: for x to be beautiful is for it to be fitting to experience delight in the presence of x . Fitting-attitude accounts are more familiar from metaethical discussions of *moral* value, but they have seen some play in aesthetics too. Franz Brentano (1906: 17), for instance, writes: 'The concept of the beautiful has to do with... [that which] elicits in us a delight with the character of fittingness.' Brentano has a complicated account of what the relevant fittingness consists in, but all that matters for us is this detail: it is *not* a matter of the delight tracking x 's independent beauty. Here too, then, the notion of beauty is not independently intelligible – independently, that is, of understanding what aesthetic delight is.

As formulated above, (*) is clearly a grounding a claim – a claim about that in virtue of which something has aesthetic value. And what it cites as ground is, as noted, a species of phenomenal consciousness. At the same time, (*) is clearly not the same kind of value-grounding claim as those considered in §§4-5. There, phenomenal consciousness grounded value in the sense that something's having a phenomenal property made it have a value property – *it* was valuable in virtue of being conscious. Here one thing is conscious (the aesthetic experience) and *something else* is valuable (whatever elicits the experience).

We may put the point as follows. In its most general form, a claim to the effect that phenomenal consciousness grounds value has the following form: for some phenomenal property P and some value property V , the fact that some x is P grounds the fact that some y is V . But this is consistent with two importantly different scenarios: one where $x = y$ and one where $x \neq y$. When $x = y$, it is natural to say that consciousness itself has the relevant value (as per our discussion in §3). In contrast, when $x \neq y$, it is no longer natural to describe the situation as one in which consciousness itself has the

relevant value. Here it is more natural to say that consciousness *enables* value, in that it is a precondition for something else having the relevant value. Of course, in virtue of enabling something having value, consciousness itself acquires an *instrumental* value (and a rather crucial instrumental value!). The fact that a painting (say) is beautiful in virtue of its eliciting an experiential response imputes *instrumental* aesthetic value on the experience – though not *intrinsic* aesthetic experience. For it is not the experience itself which is beautiful, but the painting.

I conclude that phenomenal consciousness has relatively little or no intrinsic aesthetic value, but may have an all-important kind of instrumental aesthetic value. It is worth noting, incidentally, that (*)-style subjectivist value-ground claims could in principle be made with respect to ethical and even epistemic value.¹⁷ Indeed, such claims have been made repeatedly in the history of philosophy. To that extent, the discussion in this section brings out another *general strategy* for establishing a central albeit instrumental value for consciousness.

Conclusion

The primary purpose of this paper was to sketch out a general framework within which the question of the value of consciousness could be addressed in a systematic and rigorous manner. In the first instance, the question is usefully factorized into eight types of question: about epistemic, prudential, moral, and aesthetic value, either intrinsic or instrumental. To make progress on these questions, we must first have clear views on what the intrinsic values in each domain are. Once we do have such views on the table, we can use standard philosophical techniques of analysis and argumentation to try and establish various relatively concrete theses. The more comprehensive our collection of such theses, the more comprehensive our account of the value of consciousness becomes. And the conjunction of full answers to the 48 specific questions distinguished in Part I would come close to constituting a complete account of the value of consciousness.¹⁸

References

- Allen, K. Forthcoming. 'The Value of Perception.' *Philosophy and Phenomenological Research*.
- Annas, J. 1993. *The Morality of Happiness*. Oxford and New York: Oxford University Press.
- Annas, J. 2008. 'The Phenomenology of Virtue.' *Phenomenology and the Cognitive Sciences* 7: 21-34.
- Block, N.J. 1995. 'On a Confusion About the Function of Consciousness.' *Behavioral and Brain Sciences* 18: 227-247.
- Bourget, D. 2017. 'The Role of Consciousness in Grasping and Understanding.' *Philosophy and Phenomenological Research* 95: 285-318.
- Bramble, B. 2016. 'A New Defense of Hedonism about Well-Being.' *Ergo* 3: 85-112.
- Brentano, F.C. 1889. *The Origin of Our Knowledge of Right and Wrong*. Trans. R.M. Chisholm & E.H. Schneewind. London: Routledge and Kegan Paul, 1969.
- Brentano, F.C. 1906. 'Of Beauty.' In his *Grundzuge der Ästhetik*. Bern: Francke Verlag, 1959.
- Brogaard, B. and E. Chudnoff forthcoming. 'Consciousness and Knowledge.' In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press.
- Campbell, J. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- Chalmers, D.J. 2003. 'The Content and Epistemology of Phenomenal Beliefs.' In Q. Smith and A. Jokic (eds.), *Consciousness: New Philosophical Perspectives*. Oxford and New York: Oxford University Press.
- Chudnoff, E. 2011. 'The Nature of Intuitive Justification.' *Philosophical Studies* 153: 313-333.
- Chudnoff, E. 2012. 'Presentational Phenomenology.' In S. Miguens and G. Preyer (eds.), *Consciousness and Subjectivity*. Frankfurt: Ontos Verlag.
- Chudnoff, E. 2015. *Cognitive Phenomenology*. London and New York: Routledge.
- Conee, E. 1994. 'Phenomenal Knowledge.' *Australasian Journal of Philosophy* 72: 136-150.
- Crisp, R. 2006. 'Hedonism Reconsidered.' *Philosophy and Phenomenological Research* 73: 619-645.
- Cutter, B. 2017. 'The Metaphysical Implications of the Moral Significance of Consciousness.' *Philosophical Perspectives* 31: 103-130.
- Danto, A. 2002. 'The Abuse of Beauty.' *Daedalus* 131: 35-56.

- De Brigard, F. 2010. 'If You Like It, Does It Matter If It's Real?' *Philosophical Psychology* 23: 43-57.
- van der Deijl, W. Forthcoming. 'Is Pleasure all that is Good in Experience?' *Philosophical Studies*.
- Elgin, K.Z. 'Emotion and Understanding.' In G. Brun, U. Dogluoglu and D. Kuenzle (eds.), *Epistemology and Emotions*. London: Ashgate.
- Fine, K. 2012. 'Guide to Ground.' In F. Correia and B. Schnieder (eds.), *Metaphysical Grounding*. Cambridge: Cambridge UP.
- Fletcher, G. 2013. 'A Fresh Start for the Objective List Theory of Well-Being.' *Utilitas* 52: 206-220.
- Gertler, B. 2001. 'Introspecting Phenomenal States.' *Philosophy and Phenomenological Research* 63: 305-28.
- Ghijzen, H. 2014. 'Phenomenalist Dogmatist Experientialism and the Distinctiveness Problem.' *Synthese* 191: 1549-1566.
- Giustina, A. Forthcoming. 'Introspection without Judgment.' *Erkenntnis*.
- Goldman, A.I. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.
- Gordon, E.C. 2017. 'Understanding in Epistemology.' *Internet Encyclopedia of Philosophy*.
- Hare, R.M. 1981. *Moral Thinking*. Oxford: Oxford University Press.
- Heathwood, C. 2007. 'The Reduction of Sensory Pleasure to Desire.' *Philosophical Studies* 133: 23-44.
- Heathwood, C. 2016. 'Desire-Fulfillment Theory.' In G. Fletcher (ed.), *The Routledge Handbook of Philosophy of Well-Being*. London and New York: Routledge.
- Horgan, T. and U. Kriegel 2007. 'Phenomenal Epistemology: What is Consciousness that We Should Know It so Well?' *Philosophical Issues* 17: 123-144.
- Hume, D. 1757. 'Of the Standard of Taste.' In his *Four Dissertations*. London: A. Millar in the Strand.
- Jackson. 1982. 'Epiphenomenal Qualia.' *Philosophical Quarterly* 32: 127-136.
- Kahane, G. and J. Savulescu 2009. 'Brain Damage and the Moral Significance of Consciousness.' *Journal of Medicine and Philosophy* 34: 6-26.
- Kant, I. 1797. *The Metaphysics of Morals*. Trans. M. Gregor. Cambridge: Cambridge University Press, 1996.
- Kidd, C. 2011. 'Phenomenal Consciousness with Infallible Self-Representation.' *Philosophical Studies* 152: 361-383.
- Kriegel, U. 2009. *Subjective Consciousness: A Self-Representational Theory*. Oxford: Oxford University Press.
- Kriegel, U. 2013. 'Animal Rights: A Non-Consequentialist Approach.' In K. Petrus and M. Wild (eds.), *Animal Minds and Animal Ethics*. Bielefeld: Transcript Verlag.

- Kriegel, U. 2015. *The Varieties of Consciousness*. Oxford and New York: Oxford University Press.
- Kriegel, U. 2017. 'Dignity and the Phenomenology of Recognition-Respect'. In J.J. Drummond and S. Rinofner-Kreidl (eds.), *Emotional Experience: Ethical and Social Significance*. London and New York: Rowman & Littlefield.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lee, A.Y. 'Is Consciousness Intrinsically Valuable?' *Philosophical Studies* 175: 1-17.
- Levy, N. 2014. 'The Value of Consciousness.' *Journal of Consciousness Studies* 21: 127-138.
- Levy, N. and J. Savulescu 2009. 'Moral Significance of Phenomenal Consciousness.' *Progress in Brain Research* 177: 361-370.
- Moore, G.E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Murphy, M. 2001. *Natural Law and Practical Rationality*. Oxford and New York: Oxford University Press.
- Pautz, A. 2017. 'The Significance Argument for the Irreducibility of Consciousness.' *Philosophical Perspectives* 31: 349-407.
- Peterson, A. and T. Bayne 2018. 'Post-Comatose Disorder of Consciousness.' In R. Gennaro (ed.), *Routledge Handbook of Consciousness*. London and New York: Routledge.
- Pryor, J. 2000. 'The Skeptic and the Dogmatist.' *Noûs* 34: 517-549.
- Seager, W. 'Consciousness, Value and Functionalism.' *Psyche* 7 (20).
- Shepherd, J. 2018. *Consciousness and Moral Status*. London and New York: Routledge.
- Sidgwick, H. 1874. *The Methods of Ethics*. London: Macmillan.
- Siewert, C. 1998. *The Significance of Consciousness*. Princeton NJ: Princeton University Press.
- Siewert, C. 2013. 'Speaking up for Consciousness.' In U. Kriegel (ed.), *Current Controversies in Philosophy of Mind*. London and New York: Routledge.
- Siewert, C. Forthcoming. *Subjectivity and Understanding*.
- Silverstein, M. 2000. 'In Defense of Happiness: A Response to the Experience Machine.' *Social Theory and Practice* 26: 279-300.
- Singer, P. 1975. *Animal Liberation*. New York: Random House.
- Smithies, D. 2014. 'The Phenomenal Basis of Epistemic Justification.' In J. Kallestrup and M. Sprevak (eds.), *New Waves in Philosophy of Mind*. Basingstoke: Palgrave-MacMillan.
- Strawson, G. 2011. 'Cognitive Phenomenology: Real Life.' In T. Bayne and M. Montague (eds.), *Cognitive Phenomenology*. Oxford: Oxford University Press.

¹ A positive answer to the first question clearly requires that *all* species of phenomenal consciousness have value. But it also requires more than that. For suppose that pluralism is true of ethical value, such that V_1 and V_2 are two mutually irreducible intrinsic ethical values. And suppose that species S_1 and S_2 of phenomenal consciousness exhibit V_1 while species S_3 and S_4 exhibit V_2 . Then there would be no intrinsic ethical value that phenomenal consciousness as such would exhibit – even though for each of its species there would be an intrinsic ethical value that *it* exhibits.

² This may be either because consciousness' epistemic value is more salient than its ethical and aesthetic value, or (as I rather suspect) because epistemology lies nearer the philosophy of mind than do ethics and aesthetics.

³ Some naïve realists, notably Keith Allen (forthcoming), have argued that perceptual experience, insofar as it provides us with contact with reality, is intrinsically rather than instrumentally valuable. It is not entirely clear from Allen's discussion, however, whether he has epistemic or ethical value (or both) in mind. In any case, I am skeptical of Allen's argument. The argument is that naïve realism is the only view of perceptual experience that can account for its intrinsic value as constituting contact with reality, because other views must claim that it is only in belief (and action) that contact with reality is achieved, and so perceptual experience is valuable only insofar as it is instrumental to the formation of beliefs (and actions). But this is wrong: most opponents of naïve realism would claim that contact with reality is established in veridical perceptual experience; it is just that it is the veridicality of the experience, not its sheer occurrence, that establishes contact. (An appeal to belief here would be strange, since most opponents of naïve realism think that false belief and true belief belong to a single psychological kind just as veridical and falsidical perceptual experience do.)

⁴ It might be argued that the type of introspective awareness of our conscious states that is infallible is itself conscious, so that this assigns at least one species of consciousness – introspective consciousness – a kind of intrinsic epistemic value. The problem, though, is that the fact that conscious introspections are infallible does not show that they are so *in virtue of being conscious*. (Indeed, most of the cited authors hold that they are so in virtue of some constitutive connection between introspection and the introspected.)

⁵ Siewert (2013) argues for a central role for consciousness in specifically *linguistic* understanding (see also Campbell 2002 for the connection between consciousness and linguistic reference). But this kind of understanding could be claimed to reduce to a species of knowledge: knowledge of meaning (in Campbell's case, knowledge of reference).

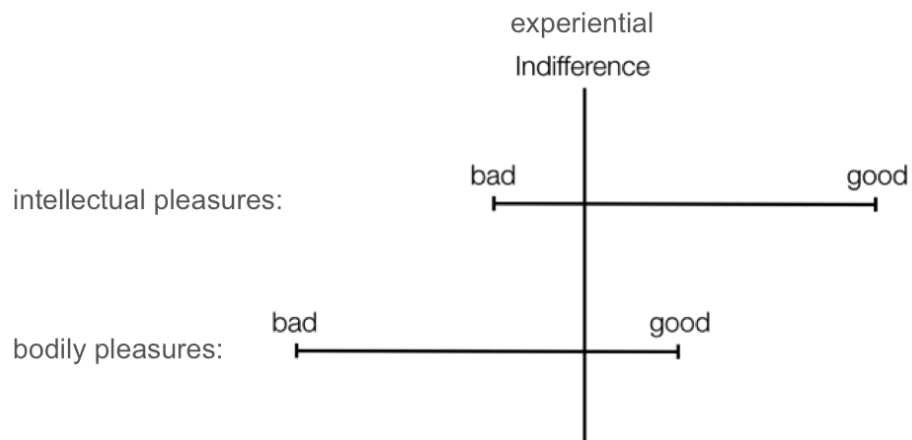
⁶ This is partially motivated by the observation that pleasures can be extremely varied phenomenologically, so that what unifies them may well be simply the fact that one desires

having them (Heathwood 2007: 26). To my mind, however, the heterogeneity of pleasure is no obstacle to its bearing prudential value. The felt qualities of seeing blue, tasting honey, and feeling velvet are extremely different as well; arguably, just as different as are the felt quality of pleasure produced by a particularly well-crafted George Eliot sentence and the felt quality of orgasm. Yet as we saw in §4, many epistemologists assign epistemic value to perceptual phenomenology, as the genus of which visual, gustatory, and tactile phenomenology are species. By the same token, we can assign prudential value to pleasure phenomenology as the genus of which many species can certainly be recognized. (This is not to say that there need not be any phenomenal commonality among pleasures; rather, I suspect the commonality is as phenomenally subtle as the commonality between all forms of perceptual phenomenology.)

⁷ Moods seem to me particularly important for well-being. When I reflect on periods in my life where my life ‘was going well,’ these periods seem to feature prominently a preponderance of good moods. Punctate pleasures played a role too, but mostly in their capacity as causes of the occurrence and maintenance of good mood. The fact that moods have a longer phenomenal half-life, and have a pervasive influence on one’s other conscious experiences, seem to give them this special role.

⁸ Joshua Shepherd (2018: 31) defends a somewhat similar view that pins prudential value on what he calls ‘affective-evaluative’ phenomenal properties – properties whereby phenomenal states are both affectively valenced and evaluate their objects. In my opinion, however, the temptation to add evaluative phenomenal properties to the mix is due entirely to the substantial correlation between evaluative and valenced properties: pleasure, contentment, hope, and so on both feel good and evaluate their objects positively, while pain, embarrassment, anger, and the like both feel bad and evaluate their objects negatively. But it is unclear why the fact that objects get evaluated through them should affect one’s own well-being. And when the correlation breaks down, our intuition is that valence rather than evaluation is what matters. Thus, yearning for *x* feels bad but evaluates *x* positively. Intuitively, other things being equal having a lot of yearning in one’s life makes life worse for one.

⁹ Various strategies for vindicating the swine intuition can be found in the extant literature, but let me float one to my knowledge entirely absent from it. (It is modeled on a structurally analogous point made by Nicolai Hartmann in a different context.) The reason the ‘intellectual’ pleasures are more valuable than the ‘bodily’ pleasures is, on this conjecture, that the point of experiential indifference lies in different places on their pleasure-pain scales. Reading a bad poem involves very little pain, whereas one good poem can fill us with euphoric exaltation and deep satisfaction; this is why we are willing to read fifty poems that do nothing for us for the prospect of reading one that shakes us.⁹ In contrast, the pleasure of eating a good strawberry is far greatly outweighed by the pain of eating a rotten strawberry. We would never be willing to eat fifty rotten strawberries for the prospect of getting the one delicious one in the box. These examples may admittedly have various distracting features, but the general theoretical point is summarized in the figure below.



¹⁰ Thanks to Géraldine Carranante for making me see this.

¹¹ Of course, we all value satisfying our desires. But it is less clear whether what we value there is (i) our desires' satisfaction condition being met, (ii) the pleasant feeling of satisfaction that attends (i) when we are aware of (i), or (iii) both.

¹² A hedonist may object that novelty, self-understanding, and so on are themselves pleasurable. But this seems to be so in a merely causal and not constitutive sense: novelty tends to *produce* pleasure, it is not *itself* a pleasure. So van der Deijl is right that his theory is not a form of hedonism. The hedonist may still suspect that it is the tight causal link to pleasure that tempts our intuition toward this kind of list-experientialism.

¹³ This is not how these views are standardly characterized in the first instance (often they are characterized rather in terms of what makes an action morally right), but standard characterizations have more or less immediate implications for what makes a life morally good. And this latter question complements nicely the question of well-being, as suggested at the opening of §5.

¹⁴ Moreover, according to Annas (2008) there is a generic phenomenology characteristic of the manifestation of virtue in general, namely, the phenomenology of *flow* – of acting in the right way in an effortless and unconflicted manner. The person who gives a homeless a dollar only after overcoming strong internal resistance, perhaps by deliberately overriding her self-interested instincts, is not as generous as the person who gives the dollar as a matter of course.

¹⁵ In the area of *animal* ethics, too, the topic of consciousness comes up time and again. In particular, the notion of 'sentience' is central, where 'sentience' is essentially the capacity to experience pain and pleasure. Traditionally, prominent defenders of animal ethics are consequentialists who insist that animal pleasure and pain should figure in the consequentialist

moral calculus (Singer 1975). In contrast, Kantian defenses of animal rights do not tend to put any notable emphasis on phenomenal consciousness. It is perfectly possible, of course, to hold that an animal's possession of phenomenal consciousness confers dignity on them and requires us to treat them as ends rather than mere means (Kriegel 2013).

¹⁶ I am grateful to Takuya Niikawa for making me see clearly this connection between consciousness, awe, and the sublime.

¹⁷ In particular, the claim that x is morally good just if x is disposed to elicit a pro attitude in the right subjects has considerable historical pedigree. And even the claim that p is true just if p is disposed to elicit a justified belief that p in the right subjects has seen its fair share of philosophical defenders. We must keep in mind, though, that for such claims to support an instrumental value for *consciousness*, they must insist that the relevant pro attitudes and beliefs be *conscious* – the occurrence of unconscious pro attitude or belief would be insufficient.

¹⁸ This paper is influenced by an immense number of philosophical exchanges over the years. Exchanges with Lorenza D'Angelo, Anna Giustina, Tricia Magalotti, Olivier Massin and especially Charles Siewert stand out in my memory, but there were many others. I have also benefitted from presenting it at the Jean Nicod Institute, and am grateful to the audience there, in particular Géraldine Carranante, Anna Giustina, Michele Impagnatiello, Slawa Loev, and Takuya Niikawa.