Iride, 20, 2007, 83-101 (in Italian).

The Fusion of Fact and Value

Catherine Z. Elgin

Harvard University

The fact and value dichotomy is a mainstay of contemporary philosophy and is thoroughly endorsed by popular opinion. Facts are supposed to be rigid and objective; they are, it is said, what they are, regardless of our attitudes toward them. Values are considered malleable and subjective; they are held to be what they are because of our attitudes toward them. 'You can't argue with the facts' nails down a point. 'That's a value judgment!' dismisses an opinion as ungrounded. But despite the consensus of the many and the wise, I shall argue that the fact/value *dichotomy* is untenable, and that the fact/value *distinction* is idle. The dichotomy posits a sharp divide between fact and value, with no bridge between them. The distinction posits a difference between fact and value, but allows for the possibility that there are bridging concepts, or that the difference is one of degree. I will argue that the dichotomy is discredited by thick concepts – concepts in which factual and evaluative factors are fused.[1] The difference proves idle, not only because no clear line can be drawn between factual and evaluative concepts, but also because so-called "factual" judgments and "value" judgments are justified in the same way and both figure crucially in the systems of thought within which we judge and justify. Because justification is holistic and involves both world guided and action guiding considerations, there is no hope of and no need to systematically distinguish

between the contributions of fact and value. Justification of any sort requires appealing to resources on both sides of the supposed divide.[2]

There is no denying that the idea of a fact/value dichotomy is intuitively plausible. When we contrast concepts like *potato* and *quark* with concepts like *right* and *good*, the difference between facts and values seems stark. Whether an item instantiates 'potato' or 'quark' seems to depend entirely on the way the mind-independent world is. To discover whether such a concept is instantiated, we employ empirical methods. Moreover, there is evidently nothing in particular we ought to do in consequence of our finding. Indifference to the instantiation of a concept like 'potato' is unobjectionable. 'Right' and 'good', however, function differently. That an act is right may in one way or another be bound up with what people think of the act. So the instantiation of such terms might not be (wholly) determined by the mind-independent world. Granted, this is controversial. But there being a controversy in the one case and not in the other is itself evidence of a difference between the two sorts of concepts. Moreover, something other than empirical investigation seems required to ascertain whether evaluative concepts apply. Finally, concepts like 'right' and 'good' are normative. If 'right' applies to an action, we should, *ceteris paribus*, engage in or approve of that action. Indifference is at least suspect, if not irresponsible.

Considerations like these make it seem that the conceptual realm bifurcates, factual concepts falling on one side of the divide, evaluative concepts on the other. But thick concepts, as Bernard Williams calls them, span the divide. Concepts like *cowardly, loyal,* and *truthful* are at once factual and evaluative. They are factual in that their correct application depends on the way the world is. They are evaluative in that their application

exhibits an assessment of their referents. Cowardly actions, for example, are actions that exhibit a distinct sort of fear. The term 'cowardly' does not apply if the agent is not appropriately fearful. Whether someone's behavior is cowardly thus depends on a matter of fact. So the concept is world guided. Its correct application depends on the way the world is. But to call the action cowardly is also to indicate moral disapprobation. With enough stage setting, it may be possible to characterize an action as cowardly without expressing disapproval of that particular action. But even in such a context, the use of the term intimates that the action is of a sort that is generally worthy of disapproval. It sounds strained to say 'Of course those sorts of actions are cowardly, but there is absolutely nothing wrong with them'. To call an action cowardly is to express the attitude that *ceteris paribus*, it ought not be done. Thick concepts then are action guiding.

Williams focuses on thick ethical concepts. But thick concepts occur in other domains as well. To call a writer a *hack*, a scientific investigation *slipshod*, an athlete a *clutch player*, a driver *reckless,* or a proof *valid* is to both to describe and to express an evaluation of the referent. The grounds for evaluation need not be moral. Evidently, thick concepts are ubiquitous.

The evaluative attitude expressed in the use of a thick concept need not be that of the speaker. An atheist can use the thick concept *sacrilegious*, even though she sees nothing wrong with the acts she so characterizes. In such cases, the negative valence of her term derives from her recognition that adherents consider such acts wrong on religious grounds. Her use is derivative from theirs. The user of a thick concept need not endorse its valence. But her use would be defective if she were unaware that the term has a particular valence, that it is not a neutral, descriptive term. She needs to recognize that

her descriptive term is value-laden, and to appreciate its evaluative role in the form of life it figures in. I plan to ignore such complications. Here all that is necessary is to recognize that some concepts that are at once descriptive and evaluative. For current purposes, it makes no difference whose evaluation a judgment involving a thick concept expresses.

Thick concepts straddle the supposed boundary between fact and value. This does not in itself discredit the dichotomy. The dichotomy would be tenable if thick concepts could be factored into purely factual and purely evaluative components. Then each component would be located on one side of the divide. If, for example,

(a) X is cowardly

were analyzed as

(b) X exhibits a distinctive type of fearfulness & (c) That type of fearfulness is blameworthy.

then (b) could be located on the factual side of the divide and (c) on the evaluative side. Rather than the concept itself falling on one side of the boundary, its fully analyzed counterpart would consist of sentences, each of which resides squarely on one side or the other. But thick concepts typically do not factor cleanly into purely factual and purely evaluative components. We seem unable to describe the sort of fearfulness that cowardly behavior manifests except by using terms like 'ignoble', 'shameful' or 'blameworthy'. That is, we seem unable to characterize the factual component of the thick concept without recourse to evaluative language. In thick concepts, factual and evaluative elements fuse. We can identify the ways such a concept is both world guided and action guiding, but we cannot identify separate components that perform the two functions. The

dichotomy is discredited.

That being so, the fact/value distinction looks harder to maintain. But things are not so simple, for thick concepts might be construed as hybrids. A mule is a cross between a donkey and a horse, but the existence of mules does not demonstrate that there is no difference between donkeys and horses. So if a thick concept is a hybrid that crosses a factual concept with an evaluative concept, its existence would not demonstrate that there is no difference between factual and evaluative concepts. The existence of thick concepts undermines the fact/value dichotomy, but by construing such concepts as hybrids, a fact/value distinction might still be retained. Thick concepts problematize the fact/value distinction, but they do not directly discredit it.

To see whether there is any point in retaining the distinction, let us set thick concepts aside and turn to thinner, putatively factual concepts. At first glance, it seems obvious that large portions of the world are as they are regardless of what we may think about them, or indeed whether we think about them at all. In one sense this is right, but not in a helpful sense. It follows from set theory that pretty much any collection of objects constitutes an extension.[3] There is an extension containing a dachshund, the number 3, and today's *New York Times*. A property is that which the members of an extension share. Since today's *Times*, the number 3, and the dachshund belong to the same extension, they have something in common. They share a property -- the property of belonging to that particular extension. Since any object belongs to a vast number of extensions, any object has a vast number of properties. Since just about any two objects belongs to some common extension, just about any two objects share a property. But these truths, being utterly general, are unhelpful. When we ask whether two items are

alike, the answer should not be trivial. The question we want to ask requires appeal to more restrictive criteria than co-membership in some extension or other. We want to know whether the items share important properties.

We thus devise category schemes to mark out the similarities and differences that matter. We draw the distinctions that fix the facts. We introduce the concept 'dog' to classify dachshunds, setters, beagles and the like, and the concept 'prime number' to classify 2, 3, 5 and the like. Although strictly speaking 3 is no more like 5 than it is like a dachshund, in ways that matter 3 is far more like 5 than it is like a dachshund.

Despite or even because of their clarity, scientific examples threaten to mislead. We are apt to think that constructing a biological taxonomy, for example, is simply a matter of supplying labels for what was already the case. Then prior to categorization, a dachshund was the same sort of thing as a setter, and a zebra was a different sort of thing from a horse. But these similarities and differences cannot simply be read off nature. In some respects a dachshund and a setter are alike; in others they are different. In some respects a horse and a zebra are alike; in others they are different. Since any two things are alike in some ways and different in others, mere likenesses and differences cannot settle matters of categorization. When we introduce the concept 'dog' to mark out what dachshunds, setters and beagles share, or 'prime' to mark out what 2, 3, and 5 share, we distinguish important from unimportant similarities. But that something is, or is not, important is a value judgment.

Such judgments are not arbitrary. Normally they are grounded in an appreciation of why a particular classificatory scheme is wanted; and this depends both on what we already believe about the subject and on our interests or goals with respect to it. If we

seek to understand heredity, for example, it is reasonable to group together animals that interbreed. Then despite their obvious differences, dachshunds and setters belong together and zebras and horses belong apart. If our goal is to understand gross anatomy, however, it might make sense to consider zebras and horses the same sort of thing.

More general considerations also play a role. If we seek a scheme suitable for science, scientific values and priorities should be respected. Membership in kinds should be determinate and epistemically accessible. As far as is feasible, there should be no ambiguity, no indeterminacy, and no irresolvable uncertainty about membership in the kinds. The classificatory system should foster the formulation and testing of fruitful generalizations and the design of illuminating models, and should perhaps build on or mesh with other scientific classifications of the same or adjacent domains. In constructing a system of categories suitable for science, we make factual judgments about what the values of science are, practical judgments about how they can be realized, and evaluative judgments about the extent to which they should be respected.

Science deploys streamlined categories in hopes of generating exceptionless, predictive, quantitative laws. Literature has different objectives. It seeks the capacity to characterize the particular, the exceptional, the unique. So schemes that suit literature are apt to exhibit different features from schemes that are appropriate for science. Scientific vices – ambiguity, imprecision, immeasurability, and indeterminacy – are often literary virtues.[4] Henry James's complex characterizations of social and emotional life require a baroque conceptual scheme whose categories intersect in subtle, intricate ways. Analogously intertwined categories may be needed to achieve the sorts of understanding that biographers, historians, psychoanalysts and self-reflective subjects seek. What

similarities are worth recognizing, and what categories are worth contriving thus depend on our goals and objectives in categorizing.

A category scheme provides the resources for recognizing certain likenesses and differences, for stating various truths and falsehoods, for exhibiting particular patterns and discrepancies, for drawing specific distinctions, for demarcating conceptual boundaries. Ordinarily, we do not construct a category scheme from scratch. Rather, we work with what we have got. We start with what we take to be the most appropriate available division of the domain, and revise it as necessary. The contours of a category scheme thus typically result from tinkering. When available categories do not suit our purposes, -- when they do not draw the distinctions that we want to draw, or when they make it hard to acknowledge similarities that we consider significant -- we draw new lines.

We make adjustments with more or less specific purposes in mind and incorporate into the scheme the values and priorities that we think will serve those purposes. But the resulting scheme may have features we do not intend. Some result from oversights. In an effort to regiment an unruly list of disabilities, we might decide to characterize as disabilities all and only significant deviations from normal human functioning. This criterion turns out to be ill-advised since not all such divergences carry negative consequences. Perfect pitch, for example, is a deviation from normal pitch perception, but it is an asset, not a liability. Other schemes contain inadvertent holdovers from prior systems. Flutes are still characterized as woodwinds even though they are no longer made of wood. Perhaps ‘woodwind’ is a misnomer, suggesting a material basis for the musical features the instruments in its extension share. Yet others schemes have

unintended byproducts. A taxonomy that distinguishes classes of animals on the basis of physiology might exclude egg layers from the class of mammals and lactaters from the classes of birds, reptiles, amphibians, and fish. This seemingly reasonable division lacks a classification for monotremes. Having been contrived before monotremes were discovered, it had no reason to provide a classification for egg-laying lactaters.

A category scheme is a system of predicates that group together items in a domain. By framing such a scheme we equip ourselves to state various facts. So long as there is a distinction between a predicate's seeming to apply and its actually applying, assertions that it does or that it does not apply are objective. Once such a scheme is in effect, the truth values of the sentences deploying it are independent of our views about them. Under our first scheme, 'Perfect pitch is a disability' is true. Under our second, 'Flutes are woodwinds' is true. Under our third, 'Platypuses are neither mammals, birds, amphibians, reptiles nor fish' is true. Once we have fixed the criteria for being a disability, a woodwind, or a mammal, there is no room for negotiation. The facts are as stated. The category schemes in question are defective, not because they fail to fit the facts, but because they are in one way or another unsuitable. The first supplies counterintuitive classifications; the second uses potentially misleading labels; the third provides no positive classification for some of the items it ought to be able to classify. These are all genuine defects. At least the first and the third are serious enough to warrant rejection of the schemes that yield them. But the objection is not that they fail to fit the facts; it is that they fit the wrong facts. Given what we want a classification of disabilities, a biological taxonomy, and (arguably) a classification of musical instruments for, the facts they yield are not the facts we should be interested in.

In devising a category scheme, we draw lines that mark out certain likenesses and differences as significant. Our interests determine what lines we should draw. Since our interests are multiple, there is no reason to expect that a single system of categories will serve them all. Botany has no need for, hence no incentive to devise, categories suitable for musicology. It thus lacks the resources for characterizing the *Jupiter Symphony* as in the key of C major. This, of course, no defect in botanical categories. They are designed with other ends in view. The independence of musicological and botanical categories does not lead to relativity, for no truth expressible in botanical terms is in tension with any truth of musicology. Independent category schemes may complement each other or be indifferent to each other.

Relativity emerges when systems clash – when the way things are according to one system is at odds with they way they are according to another. This might occur when schemes agree on necessary conditions for the application of a shared term, but diverge over sufficient ones, or when they agree about clear instances of a term but diverge over its applicability in borderline cases, or when they diverge about the classification of disputed items in the domain.

According to a familiar scheme for demarcating zoological classes, mammals are lactating creatures who give birth to live progeny and birds are non-lactating creatures who lay eggs. In the vast majority of cases, these criteria work well in distinguishing mammals from birds. But the platypus poses a problem, being an egg-laying, lactating animal. To accommodate the platypus, the scheme must be revised. Three obvious revisions are: (a) to introduce another class that comprehends lactating egg-layers; (b) to locate platypuses in the class of mammals, in effect decreeing that how neonates are

nurtured is, for classificatory purposes, more important than how they are born; and (c) to locate platypuses in the class of birds, in effect making the opposite assessment. Although they disagree about how platypuses ought to be classified, all three options seem reasonable. It is hopeless to attempt to decide among them by appealing to the facts, for the three schemes agree on all the relevant biological facts. They diverge over how best to accommodate those facts. All three yield acceptable accommodations. Relativity results. Relative to (a) and (c), 'The platypus is a mammal' is false; relative to (b) it is true. Each scheme yields a determinate, objective answer to the question 'Is the platypus a mammal?' Once we have chosen a scheme, the matter is settled. But the antecedent zoological facts provide no basis for choosing one of the schemes over the others.

Relativity also emerges when schemes that agree on necessary conditions diverge over sufficient ones or agree on clear cases but diverge over which borderline cases are to be included in the extension of a term. One scheme might call any unit in a multiple-family dwelling 'an apartment'. Another might restrict the term to rental units. So whether a denizen of a condominium lives in an apartment is decided differently under the two schemes. Relative to the first, 'Pat lives in an apartment' is true; relative to the second, it is false.

It might be urged that such schemes do not really clash. What the apparent clashes show is that terms are ambiguous. 'Mammal' has one extension in (a), another in (b), a third in (c). 'Apartment' has a wide extension under the first scheme and a narrower extension under the second. The diagnosis strikes me as correct. The terms in question are evidently ambiguous, since their extension varies with the scheme they

belong to. But this does not show that the schemes do not clash. Rather it shows that the clash is not a dispute over a matter of fact. It is a disagreement over the best way to characterize the domain. Each scheme highlights certain features of the items in its domain and obscures others. To decide which scheme is preferable is to decide which features ought to be highlighted and which ones it is acceptable to overshadow. Once we have chosen (b), platypuses are, as a matter of objective fact, mammals. There is no disputing that fact, even though they could just as well have been classified as birds.

Such objectivity may seem spurious if we devise category schemes to reflect our interests and can switch schemes at will. If a sentence that is true according to one scheme is false according to another, why can't we just choose a scheme whose facts we like? Things are not so simple, for rightness requires more than the capacity to state facts. To be correct, a judgment needs to state the right facts – ones that are relevant, apt, and suited to our purposes. So we need to deploy appropriate categories. For biological purposes, we can correctly classify an individual as a person on the basis of his DNA. Such a classification will not do for moral, political or what Locke called 'forensic' purposes, though[5]. There we want a category that marks out those individuals who have the capacity to act responsibly and to participate in public life. Not everyone with human DNA can do that. We would be wrong to deploy the biological scheme in deciding who is eligible to vote, because the class of persons it recognizes is not a class that suits our purpose. The facts it enables us to state are genuine facts, but they are the wrong facts.

We cannot, in any case, construct whatever we please. Our constructions are limited by our ingenuity, our resources, and the laws of nature. Over the course of history, we have increased the efficiency of our machinery enormously, and no doubt will

continue to do so. But we cannot construct a perpetual motion machine. We have devised computer programs to calculate the value of π to millions of decimal places. But we will never construct a machine capable of calculating the last digit in the decimal expansion of π. Even more limited aspirations are unrealizable. It would be nice to develop a recipe for a delicious, non-fattening tiramisu. But as things stand, replacing sugar, cream and mascarpone with non-fattening substitutes destroys the flavor. Some of the limitations are irremediable; others will eventually be overcome. Perhaps some day we will solve the tiramisu problem, but if the second law of thermodynamics is correct, we will never devise a perfectly efficient engine. Nor if mathematics is correct, is there any hope of computing the final digit in the decimal expansion of π; for there is no such digit. Construction is something we do, and we cannot do whatever we wish. Our capacities are limited, and our aspirations are often jointly unrealizable. This is so for category schemes just as it is for other constructions. So even though we construct the categories that fix the facts, we cannot hope to construct categories capable of converting every fantasy into fact.

Some separately realizable desiderata are jointly unrealizable. It might be desirable for a scheme to consist of precise categories that are easily applied. Unfortunately, the more precise the categories, the harder they are to apply. It is, for example, much easier to tell whether a ball is black or white than it is to tell whether a ball is ecru or ivory. The construction of a category scheme is informed not just by first-order desiderata, but also by second-order assessments of those desiderata. These determine what tradeoffs we are willing to make when we discover that we cannot construct a scheme that realizes all our objectives. Different tradeoffs may be reasonable.

One acceptable scheme might sacrifice precision for ease of application. A different, equally acceptable scheme, might introduce more precise categories that are harder to apply. Relative to the cruder system, two items count as the same color; relative to the more refined one, they are different colors. The design of a descriptive apparatus then is imbued with value judgments, judgments of the form that it is better to do things this way rather than that, or that it is more important to realize these objectives than those, or that this revision is permissible, but that one sacrifices too much. Any statement or judgment of fact, being cast in terms of categories we construct, is informed by judgments of value.

So far, I have been speaking as though the structure of a category scheme is determined by antecedently fixed, sharply defined values and priorities. Factual judgments have been characterized as constrained by such schemes, not as constraints on them. This is an oversimplification. Among the considerations that figure strongly in the design of a category scheme are convictions about the matters of fact it ought to reflect. We would, for example, be extremely reluctant to countenance a color scheme that classified ripe tomatoes as blue. Firm, uncontroversial antecedent judgments of fact provide constraints on classification.

Reality may resist schematizations we would like to impose. It is not unreasonable to want a system that enables us to simultaneously specify the position and momentum of an electron. But this cannot be done. We can determine an electron's position or we can determine its momentum, but it is impossible to simultaneously determine both. So the contention that schemes impose order on the world cannot be construed as claiming that the world can be ordered however we like.

Desiderata may initially be inchoate or incomplete. They are refined, elaborated,

and extended through feedback from the facts. We might start out with the vague goal of constructing some quantitative scheme for colors. We discover that by classifying on the basis of wavelengths, we can incorporate the scheme for color into a broader scheme that also comprehends electromagnetic radiation beyond the visible spectrum. Even though it was not among our initial desiderata, the fit with the broader scheme is an attractive feature. Indeed, it is so attractive that it is likely to be included as a desideratum in subsequent revisions of the system.

A category scheme is an element of a system of thought and is constructed in tandem with the rest of the system. We begin with a collection of considerations that inform our theorizing. It includes putatively factual judgments, goals, constraints on the realization of our goals, as well as first- and second-order values and priorities. These considerations are apt to be incomplete and in tension with one another. Systematizing involves articulating, extending, revising and rejecting considerations to generate accord. The goal is a system of considered judgments in reflective equilibrium. Its components should be reasonable in light of one another, and the system as a whole should be reasonable in light of our antecedent relevant commitments, when assessed by currently acceptable standards of relevance and reasonableness.[6]

There is no requirement that every element of a tenable system must be separately secured by antecedent commitments. Sometimes it is reasonable to integrate into a system a consideration that has no independent support. There is, for example, no direct evidence for the existence of positrons. So considered in isolation, there would be no reason to believe that positrons exist. But physics does not and should not consider the matter in isolation. Ontological commitment to positrons derives from their role in

physical theory. Physics commits itself to the existence of positrons on the basis of the well-founded symmetry principle that every particle has an anti-particle. An atomic theory that preserves symmetry by recognizing a particle for which there is no direct evidence is deemed by physics to be more reasonable than the one violates symmetry by refusing to admit positrons in the absence of direct evidence of them. The acceptability of the theory it belongs to thus confers acceptability on the claim that positrons exist.

Sometimes it is even reasonable to violate antecedent commitments. The theory of relativity conflicts with antecedent opinions about simultaneity by contending that, contrary to everyday appearances, events that are simultaneous relative to one frame of reference may be successive relative to another. Again the explanation lies in the contribution of the commitment to the overall theory. Accepting a counterintuitive contention about simultaneity yields a more reasonable overall astrophysical theory than abiding to the dictates of 'common sense' does. Radical departures from some antecedent commitments sometimes yield a system that is best on balance.

Reflective equilibrium requires more than mere coherence. A system is coherent if its components form a mutually supportive structure. Since each is reasonable in light of the others, the system is in equilibrium. But a coherent system can be utterly implausible. A good nineteenth century novel, for example, is coherent. But we would, and should, be disinclined to take it as face value, for its internal coherence affords no evidence for its truth. Reflective equilibrium is an equilibrium that on reflection we can accept. Our antecedent commitments constitute the basis for reflection. Because they are our prior best guesses about how things stand, a system of thought that improves on them has a claim to our epistemic allegiance. But we can only judge the system to be an

improvement if we assess it against those commitments. That a system is as reasonable as any available alternative in light of our antecedent commitments thus affords reason to accept it. Such a system need not incorporate the antecedent commitments it is judged against, but it should show why, with hindsight, they seemed as reasonable as they did. On this score, the theory of relativity fares well. Because the relativity of simultaneity is not evident at slow speeds and short distances, it is not surprising that our forebears, who only had access to (in astronomical terms) slow speeds and short distances thought simultaneity was absolute.

Deferring to antecedent convictions might seem to be an invitation to entrench old errors. It is not. For considered judgments are subject to revision or revocation in the course of theorizing. They deserve some deference, though, because they are our current best guesses about the subject at hand. They thus possess a measure of epistemic inertia. We need a reason to give them up. Failure of commitments to mesh often provides ample reason.

The picture that emerges is holistic. The basic unit of acceptance is the theory or system of thought. Acceptability of individual components is derivative, stemming from their contribution to an acceptable system. Values infuse every system of thought. For a system reflects views about what range of antecedent commitments it *ought to* answer to, what tradeoffs are *permissible* if commitments clash, how closely it *ought to* accord with the antecedent commitments that constitute its tether. It also reflects views about *good* methods for finding about the facts it concerns and for assessing findings. These reflect commitments about second-order matters such as the *trustworthiness* of methods and the *truthfulness* of informants.

Although the values that infuse the systems we construct and accept are subject to dispute, this does not make them or the systems they figure in subjective. Quite the contrary. To say that something is subject to dispute is to say that reasons can be given for or against it. We might, for example disagree about whether a sharp-edged, highly technical vocabulary is preferable to a looser vocabulary that hews closely to ordinary language. Such a dispute turns on the relative importance of familiarity and precision, which depends heavily on what the terminology is for, who will use it, what sort of training its users are likely to have, and the downside risks of each alternative. Each side can adduce reasons. Nonetheless, the considerations favoring one side may be decisive. If the only people who are apt to use the terminology are highly trained neurologists who need to make fine distinctions with precision, the highly technical vocabulary is preferable. If the system's goal is to facilitate public discussion about scientific issues, trading off precision for widespread intelligibility is reasonable. Two points emerge: Value judgments are not beyond dispute. And disagreements about values can often be settled by appeal to considerations that all parties to the dispute can accept.[7]

I have argued that holistic considerations vitiate any attempt to construe value judgments as more subjective, or less justifiable than judgments of fact. I have focused mainly on examples from the natural sciences, for these are likely to be value-free if any judgments are. Even scientific theories, I have argued, embed and presuppose judgments of value. For this reason, I consider the fact/value distinction idle. Factual judgments must be made in contexts informed by values and value judgments must be made in contexts informed by assumptions about facts. But aside from a brief mention of truthfulness and trustworthiness, my discussion has done little to directly support the idea

that judgments about moral values are on a par with judgments of fact. Typically, support for the fact/value distinction comes from the conviction that moral claims are in some important sense different from scientific (or more generally descriptive) ones. So the possibility remains that even if a general fact/value distinction is not viable, a moral/descriptive distinction is. To address that question, we need to consider whether the account offered so far comprehends the ethical domain.

At least three lines of argument have been adduced to support the contention that ethical judgments are different from factual judgments. First, there are widespread, apparently intractable disagreements about ethical matters, whereas disagreements about scientific matters seem neither so widespread nor so resistant to resolution. Second, ethical rightness seems relative to ways of life whereas factual rightness does not. Third, if ethical judgments are objective, they state ethical facts, and ethical facts would be queer facts. Let us consider these in turn.

It is undeniable that there are longstanding ethical disagreements. But we should not too quickly conclude that they show that the ethical problems they concern cannot be objectively solved, or even that they have not already been objectively solved. In certain cases it seems, disagreement persists only because some parties to the dispute are incapable of recognizing that it has already been solved. Slavery, for example, still exists in some parts of the world. Those who enslave others evidently think that they are entitled to do so. They are simply wrong, even if they are too benighted or bone headed to see that they are. We should not take the continued existence of slavery or the sincere protestations of entitlement by enslavers as any reason to believe that 'Slavery is wrong' is subjective, or controversial, or just a matter of opinion.

There are parallels in science. Findings that defy common sense or threaten complacent assumptions are not universally accepted. So there is disagreement about, for example, whether humans evolved, and whether simultaneity is relative. Many people would answer 'no' to both questions. We do not take this to impugn the objectivity of science. Rather we recognize that some people are ignorant of the relevant facts, or are incapable of understanding or accepting them.

In other cases, however, disagreements afford prima facie evidence that a problem has not been solved. That people disagree about the permissibility of permitting a young child to donate bone marrow to a sibling may be evidence that the problem has not yet been solved. But this does not show that it cannot be solved. Every field has outstanding problems; it should be no surprise that ethics does too. If our ignorance of the cause of an ALS does not impugn the objectivity of immunology, our inability to reach consensus on a solution to the problem of child-to-child bone marrow donations should not impugn the objectivity of ethics. Both problems merely demonstrate that more work needs to be done.

Still, it should be conceded, that some ethical problems probably cannot be solved. If they are ill formed, too delicate, or just too hard, they will permanently resist solution. All of these have parallels in science. We cannot calculate the last digit of π, or count the number of stars, or exactly solve the three-body problem. But these incapacities do not undermine the objectivity of science. Neither the existence of longstanding disagreement nor the concession that we cannot solve every problem we can frame then automatically discredits the claim of a discipline to be objective.

The second line of objection is that unlike matters of fact, ethics is keyed to ways

of life. This seems plausible if we contrast ethics with physics, but far less so when we remember that the realm of the factual extends beyond physics.

In any case, not all ethical principles are keyed to ways of life. The categorical imperative and the principle of utility are universal in scope. Perhaps their applications depend on ways of life, but the principles themselves do not. But even if basic moral principles are utterly general, some specific ethical obligations seem to depend on social arrangements. Societies differ over the strength and extensiveness of family obligations, and one seems to incur the family obligations defined by whatever society one belongs to. A Korean daughter could not reasonably expect to be exempted from her obligations to her parents on the grounds that American daughters have no such obligations.

Even if this is so, it does not undermine the objectivity of ethics. Many ethical obligations arise out of contingencies. Joe is obliged to pay Mary five dollars because he borrowed the money from her. Bill is under no such obligation, because he did not borrow money. Members of a mutual aid society are obliged to support each other in ways that they are not obliged to support non-members. Both the nature and extent of the required support vary with the (explicit or tacit) agreements the members of the society have made with one another. The obligations are contingent on the agreements, but given the agreements in effect, the obligations are determinate.

Families, communities, nations, and societies might seem different, since they are involuntary associations. But there is no reason to think that this difference undermines objectivity. The obligations that arise from fundamental, involuntary social practices and institutions are interwoven into the fabric of social life.[8] Specific obligations differ because social arrangements do. Since the network of social arrangements is intricate,

there are complicated relations of dependence, interdependence, and independence, connecting obligations, expectations and other social facts. Societies differ not only over whether one has special obligations to members of one's extended family, but also over who belongs to one's extended family. Distinguishing between the austere facts of physics and the cross-culturally variable list of familial obligations makes the difference look stark. Attempting to distinguish between the cross-culturally variable list of family obligations and the other anthropologically relevant, cross-culturally variable aspects of kinship requires a line that may be impossible to draw.

Again there are parallels in the natural sciences. Whether a substance is toxic depends not only on its chemical properties, but also on the organism affected by it. A substance that is toxic to one organism is non-toxic to another. So being toxic is keyed to the target organism. It is not an intrinsic property of the substance. Still, that substance *a* is toxic to organism *b* is a brute matter of fact. Chocolate is toxic to dogs but not to humans. Peanuts are toxic to some humans but not to others. But whether chocolate or peanuts are toxic to a given organism is a determinate, objective fact.

The final objection is John Mackie's 'argument from queerness'. Moral facts, he believes, would be 'utterly different from anything else in the universe', and would require a special moral faculty to discern them. Because he can find no place for them or our knowledge of them in the material world, he concludes that there are no such facts. Moral claims are, he contends are subjective.[9] Actually, the 'argument from queerness' is not an argument, but a bald assertion.[10] Nevertheless, it has a certain prima facie plausibility, so it is worth rebutting.

Mackie's position rests on two assumptions that we need not and should not grant.

The first is a commitment to reductive materialism. Facts deemed are queer if they do not reduce to or straightforwardly supervene on material facts. The second is a commitment to a sort of flat-footed empiricism that grounds knowledge of unqueer facts in sense perception of the material world. Facts are deemed queer if sense perception does not straghtforwardly afford epistemic access to them. A wide range of facts violate these conditions. By Mackie's standards, mathematical facts turn out to be queer. They are neither reducible to nor supervenient on material facts, and our knowledge of them does not derive from sense perception. Nonetheless, we are not inclined to say that mathematical truths are subjective. Historical facts seem queer as well. No doubt the signing of the Declaration of Independence was a physical event. But neither the event nor its causes and consequences wholly reduce to or straightforwardly supervene on facts describable in purely physicalistic terms. The appearance of the ink on the parchment can perhaps be described physicalistically. What cannot be so described is that the members of a legislative body, acting on behalf of their constituents, inscribed their names, thereby declaring that the political entity they represented no longer owed allegiance to another political entity. That requires locating the signing in a network of irreducibly political and social facts – facts about conventions, rules, agreements, procedures, precedents, expectations, fears, and so forth.

The alleged queerness of such facts makes no difference. They are integral components of acceptable theories or systems of thought. If statements of queer facts were removed from such theories or assigned a merely expressive status, significantly weaker theories would result. If, for example, the only thing to be said in favor of truthfulness in scientific communication were that scientists approve of it, the confidence

placed in reports published in respected scientific journals would be misplaced. Confidence in truthful reports would be no more reasonable than confidence in reports printed in a popular font. Earlier I said that a system of thought is acceptable only if it is at least as reasonable as available alternatives. I am now arguing that systems that recognize the reality of the sorts of facts that Mackie deems queer are more reasonable than systems that exclude them, or assign them a merely expressive status on account of their queerness. That a report is truthful is a reason to believe it, not merely a likeable feature of it.

A final worry remains. I said earlier that reflective equilibrium results from a process that allows for the revision or rejection of commitments in the interests of achieving a theory that on balance we can accept. But not everything is negotiable. Torturing the innocent is wrong, even if a system that permitted it would be better on balance than systems that forbade it. The worry is misplaced. For the commitments that serve as the basis for systematization are weighted. So not all are equally susceptible to revision or rejection. If a commitment has sufficient weight, then it is effectively non-negotiable. In that case, to achieve equilibrium we must adjust other commitments to accommodate it, rather than adjusting it to accommodate other commitments. There is no possibility that a system that permitted torturing the innocents would be in reflective equilibrium, for we would not on reflection accept any such system. No commitments are absolute in the sense of being *in principle* immune to revision or rejection. But some commitments are so firmly and confidently held, and so deeply entrenched in our understanding that we have no reason to think that we will ever have grounds for giving them up.

I have argued that factual and evaluative considerations intertwine to constitute our systems of thought. Their acceptability derives from the acceptability of the systems they figure in. Because acceptability of a system is a matter of reflective equilibrium, an acceptable system must contain both factual and evaluative elements. Since both are assessed in the same way, factual judgments are not objective unless value judgments are; and value judgments are not relative unless factual judgments are. I have urged that judgments of both sorts are objective and relative. There is no sharp line between the factual and the evaluative. Thick concepts are inextricably both. Rather than construe them as hybrids, it is preferable to recognize a conceptual continuum. At the one extreme are the evaluatively austere concepts like 'potato' and 'quark'. At the other are highly value-laden concepts like 'right', 'good'. In between are not only concepts like 'cowardly' and 'truthful' but also concepts like 'valid' and 'verified'. Not even physics or mathematics is free of evaluatively thick concepts. The stereotype of factual knowledge as consisting of value free theories about the way the world is, which may subsequently be over overlaid with subjectively evaluations, is implausible. Since values infuse our lives, it is not surprising that they also infuse our understanding of our lives, our world and our place in the world.

Catherine Z. Elgin

Harvard University

[1] Bernard Williams. *Ethics and the Limits of Philosophy* (Cambridge: Harvard, 1985), 140-141.

[2] See my 'The Relativity of Fact and the Objectivity of Value,' in *Between the Absolute and the Arbitrary* (Ithaca: Cornell University Press, 1997), 176-191.

[3] The caveat 'pretty much' is required because membership in sets is limited to avoid the set-theoretical paradoxes and worries about large cardinals. These complications are irrelevant to my point.

[4] See Israel Scheffler. *Beyond the Letter* (London: Routledge, 1979), 7.

[5] John Locke. *An Essay Concerning Human Understanding,* volume 1 (Oxford: Clarendon, 1894), 445-470.

[6] See my *Considered Judgment* (Princeton: Princeton, 1996), where this approach is developed in detail.

[7] See T. M. Scanlon. *What We Owe to Each Other* (Cambridge: Harvard, 1998), 4.

[8] To say that some obligations arise from social arrangements we are born into is not to say that any putative obligation endorsed by such social arrangements is a genuine obligation. Clearly that is not so. The picture is rather that universal obligations are realized differently in different social settings. The general obligation to help others in need gives rise to specific obligations to do this or that, depending on who is in need and how they can be helped in different social environments.

[9] See John Mackie. *Ethics: Inventing Right and Wrong* (London: Penguin, 1978), 38-42. Quotation from p. 38. Mackie focuses on ethical objects such as values rather than on ethical facts. But his skepticism about ethical objects easily extends to ethical facts.

[10] I owe this point to Judith Jarvis Thomson.