

begs the question of what is driving moral decision processes. The second question we are left with is: What, if any, role would this continuum play in the moral actions of the individual? Do we use the same type of default system when we are making socially ambiguous decisions that might directly affect us and/or those around us? (For example, are there pre-judged lines that we will not cross?)

We have made the argument (Kang & Glassman 2010) that moral thought, including the type of moral judgments Knobe describes, is actually a form of what Bourdieu (1986) refers to as *cultural capital*, while moral action is a form of *social capital* (Portes 1998). The motivation and goal (Glassman 1996) of cultural capital is to signal to those around you that you are a member in good standing of the social group. It is a short-hand for the types of social interactions that allow individuals to establish affiliation through community standards. Moral judgments are one of the easiest forms of cultural capital to use to establish group membership, whether it is gossip around a community pool or the establishment of a common enemy, villain, and/or scapegoat.

We suggest that the primary moral judgment that Knobe describes is made in the service of cultural capital, and that it is more about signaling and establishing membership in a given community than “about controllability, about recency, about statistical frequency” (sect. 5.1, para. 5). The default position of what is acceptable for the businessman in Knobe’s scenario example would change dramatically based on whether you were trying to signal membership and affiliation in the Chamber of Commerce or in the Sierra Club.

The setting of a continuum and establishing of a default is a form of cultural capital, and if we are on target in our thesis (Kang & Glassman 2010), it would all but disappear when individuals are engaged in collective moral action. In moral action, individuals are less concerned with establishing a signal/symbol system for long-term group maintenance and belonging than in coming together as a group to solve a critical problem. In moral action, the focus is almost completely on the problem at hand, rather than on who should be included (and excluded) from the working group. The action is integrated with the specifics of the problem to be solved, and as the common problem dissipates, so, too, does the motivation behind the group (Putnam 2001).

We see two reasons why there is little to be gained by using the primary moral judgments of generalized situations in moral action. The first is that group membership is malleable in problem solving, and placement in the group is dependent on abilities. The second reason is that problems are dynamic and shifting, and individuals who are taking action might have to continuously abandon or change their default point based on circumstances. To take a crude example, a person with a specific default position on sharing of community resources might take a very different view if he or she is placed in charge of such resources. (For example, how would the individuals in Knobe’s academic example change if they found themselves being denied access to pens when they needed them? Or if their salary were dependent on maintaining a supply of pens?)

We take the real-world example of the recent British Petroleum (BP) oil spill to illustrate our point, similar to Knobe’s businessman who does not care about risks to the environment. Suppose, before the spill occurred, people were asked about the intentions of the president of BP if he said the company could engage in deep water drilling without harming the environment: There would be a wide array of responses, directly based on the primary judgments Knobe discussed, but judgments used to signal community belonging. If you asked an officer in an environmental group, he or she might have set the default point for acceptable action so that the greater part of the continuum led to intuitions of morally bad intentions (e.g., being willing to drill at all, or not actively investing in sustainable energy). If you asked a politician from the Gulf Region, he or she might have set a default position with far more of the continuum

devoted to a neutral position (e.g., drilling could occur as long as there were minimal safety precautions) in order to signal kinship with the oil-dependent community.

After the spill, many members of the two groups have acted together in attempting to stop the spill and reclaim the Gulf. Intuitions about good or bad intentions and the moral judgments that led to them have become secondary or even irrelevant for many working in this group, and it is considered bad form to bring them up. Ties have been established based on the need to solve the immediate problem. Once the problem has diminished, or retreats into the background, the social group will dissipate and moral judgments as cultural capital will move to the fore again. It represents a cycle of moral thought as cultural capital and moral action as social capital.

Are mental states assessed relative to what most people “should” or “would” think? Prescriptive and descriptive components of expected attitudes

doi:10.1017/S0140525X10001792

Tamar A. Kreps and Benoît Monin

Graduate School of Business, Stanford University, Stanford, CA 94305.

Kreps_Tamar@gsb.stanford.edu

monin@stanford.edu <http://www.stanford.edu/people/monin>


Abstract: For Knobe, observers evaluate mental states by comparing agents’ statements with “defaults,” the attitudes they are expected to hold. In our analysis, Knobe’s model relies primarily on what agents *should* think, and little on expectancies of what they *would* think. We show the importance and complexity of including descriptive and prescriptive norms if one is to take expectancies seriously.

If you claimed at a dinner party to have no opinion about child abuse, you would get funny looks. In Knobe’s analysis, because you should strongly oppose abuse, neutrality is tantamount to support. Similarly, expressing neutrality about women’s suffrage, which our society supports, would appear sexist. Thus, observers do not take agents’ claims at face value, but instead assess them relative to what Knobe calls a “default.” Observers essentially convert an agent’s claim to their own metric, much like converting Celsius to Fahrenheit, based on the object of judgment (e.g., helping vs. hurting the environment) and the associated “default” attitude (see our Fig. 1, Panel A).

This “default,” defined in the target article as “what sort of attitude an agent could be expected to hold toward” an object (sect. 5.2, last para.), and elsewhere (Pettit & Knobe 2009) as what any reasonable person “would” (p. 597) or “should” (p. 598) think, is thus a central part of Knobe’s model. In this commentary, we aim to analyze and clarify this concept, which we believe is more complex than Knobe lets on. There is much to be gained from such analysis, especially from distinguishing the *should* and *would* aspects of default expectations.

What influences people’s expectations about how others behave and think? Certainly, one factor, as Knobe points out, is *personal moral judgment*: we expect people to behave in (what we ourselves believe is) a moral fashion. However, two other social factors seem at least as important as personal moral judgment in determining defaults: *prescriptive norms* (how we think the *group* believes people should act) and *descriptive norms* (how we think group members *actually* act, regardless of how they should). Personal moral judgments do not always correspond to group prescriptive norms, and the default expectation often depends on the latter, as when an agnostic, hearing an American presidential candidate publicly espousing agnosticism, sees this as a forceful anti-religion stance given American norms, even if it accords with his own views. Similarly, a

Panel A



Panel B




Figure 1 (Kreps & Monin). Converting expressed attitudes into inferred attitudes by reference to a default (Panel A) or to known pressure (Panel B).

default based on descriptive norms explains why, even if I know that I (personal moral judgment) and my colleagues (prescriptive norm) believe it is better to take public transportation than to drive to work, my assessment of a colleague who drives (and whether that means she “supports” public transportation) still depends on whether I know my colleagues generally drive or not.

These examples illustrate that we evaluate other people’s choices not just relative to the default of our own personal values (how they *should* act), but also relative to what we can reasonably expect from others given our knowledge of the world (how they *would* act). Knobe privileges the *should* aspect: For example, one version of the pen/professor study (sect. 3.4) pits moral judgment against descriptive norms, and the relative importance of moral judgment is taken to support the model. Although Knobe’s issues (the environment; reasonable rules about pens) are fairly prescriptively consensual, perceived prescriptive norms could be divorced from personal moral judgment, in which case Knobe would still favor the latter. Imagine I feel strongly that eating meat is immoral, while realizing my view is the minority one. Knobe would say that I think others who express indifference are really in favor, because my *should* default is strong opposition, even though I would not reasonably expect a random stranger to share my view (*would* default).

While Knobe may be right that *should* factors matter in many circumstances, other evidence suggests the importance of *would* factors in evaluating attitudes. For example, people use prescriptive norms to infer situational pressure and correct expectations accordingly. Observers assuming strong pressures against expressing support for harming the environment can sensibly infer a suppressed *pro*-harm attitude behind an expression of indifference (Fig. 1, panel B). Similarly, a speaker advocating immediate action by a corporation to reduce pollution is perceived as more *anti*-environment when speaking to a *pro*-environment audience, where such a message is expected, than to a *pro*-business audience (Eagly et al. 1978). Here, assumed audience pressure changes the default, although participants’ personal moral judgment presumably remains constant. Ironically, this is exactly the “augmentation” process described in Kelley’s (1971) attribution theory, which Knobe dismisses as a wrongheaded “person-as-scientist” theory.

Another example where *would* matters is the impact of inter-group perceptions. A devout Catholic claiming no particular opinion on Roe v. Wade might seem more in favor than a staunch feminist making the same claim. Biernat (2005) showed that expectations associated with different groups lead to such contrast effects. (Intriguingly, Biernat’s research also suggests an *assimilation* effect with more objective measures – the

Catholic would still seem less likely to get an abortion – suggesting that Knobe’s might have found a different pattern using objective outcome measures.) Thus, group-specific descriptive norms evoked by agents’ identities influence the default.

The value of distinguishing *should* and *would* influences on defaults is further suggested by research indicating possible interactions between them. For example, personal moral judgments affect perceived norms: Research on naïve realism and social projection (e.g., Ross & Ward 1996) shows that individuals generally believe their own judgments are rational, objective, and ethically appropriate, and therefore overestimate the similarity of others’ attitudes. Also, norms can influence personal judgment: People’s desire to fit in can lead them to change their own judgment to conform to perceived norms (e.g., Asch 1956). Further, descriptive norms are sometimes inferred from prescriptive norms, and vice versa (Prentice & Miller 1996).

In summary, we believe Knobe’s model makes a valuable addition to our understanding of defaults and social judgments, but it seems to be unreasonably limited to factors based on “should”; for a fuller understanding of what determines people’s default expectations, the model could be enriched by including other factors based on “would,” such as group descriptive and prescriptive norms. Including these factors – which often have little to do with morality – might dilute the model’s focus on how moral considerations suffice social judgment, but such a change seems warranted given the important role of non-moral factors in determining default expectations. We hope future research will extend Knobe’s model to include such factors.

Understanding the adult moralist requires first understanding the child scientist

doi:10.1017/10.1017/S0140525X10002037

Tamar Kushnir and Nadia Chernyak

Department of Human Development, Cornell University, Ithaca, NY 14853.

tk396@cornell.edu nc98@cornell.edu

http://www.human.cornell.edu/che/bio.cfm?netid=tk397

Abstract: Children learn *from* people and *about* people simultaneously; that is, children consider evidentiary qualities of human actions which cross traditional domain boundaries. We propose that Knobe’s moral asymmetries are a natural consequence of this learning process: the way “child scientists” gather evidence for causation, intention, and morality through early social experiences.

Knobe’s “person as moralist” view contests two related claims about human cognition: that it is clustered by discipline, much as university departments are, and that cognition in two “scientific” disciplines – folk psychology and causal inference – is analogous to scientific inquiry. Knobe then presents evidence that the psychology of intention and causation are “suffused with moral considerations” (sect. 5.3, para. 3), by which he means to show that there is neither a separation between disciplines, nor can reasoning about scientific topics be considered “scientific.”

We suggest another perspective on these moral asymmetries: that they are, at least in part, the consequences of early links between causal learning and social learning. Specifically, they are the result of how, as children, we gather evidence for such learning by observing and interacting with people. The adult moralist recruits knowledge gained from years of social evidence gathering – years spent learning *from* people and *about* people simultaneously. Therefore, to understand the adult moralist we must first understand her predecessor – the child scientist.

For a long time, developmental psychologists studied children’s knowledge separately, according to domain. Some research examined early causal reasoning – intuitions about spatio-temporal relations (Leslie & Keeble 1987; Oakes &

Cohen 1990), causal mechanisms (Bullock et al. 1982; Schulz 1982), and the use of statistical cues in causal judgments (Gopnik et al. 2001; Sobel & Kirkham 2006). Other research focused on children’s “mind-reading” abilities – what they knew about the intentions, desires, beliefs, and knowledge states underlying human actions (e.g., Lutz & Keil 2002; Repacholi & Gopnik 1997; Wellman 1990; Woodward 1998). Others sought to understand children’s knowledge of social categories (Bigler & Liben 2007; Heyman & Gelman 2000), and still others focused on developing moral and conventional knowledge (e.g., Turiel 1983). The picture that emerged from these separate subfields is a lot like the mental university described by Knobe – separate departments for separate knowledge structures.

The domain-specific approach has led to important discoveries about the content of early physical, biological, psychological, and social and moral knowledge. However, trying to apply this approach wholesale to learning processes has been less fruitful. Take causal learning: spatio-temporal cues and mechanism knowledge are useful, but are often unavailable. Statistical cues are also useful, but cannot help distinguish between causes and spurious correlations. Most often, ordinary causal learning depends on social interaction; evidence comes from doing things and watching others do things. Human actions are a child scientist’s natural causal experiments (Gopnik et al. 2004; Schulz et al. 2007).

Importantly, along with physical evidence (e.g., toys making noise, milk spilling, sticks breaking), causal actions contain valuable social evidence (a knowing glance at the right button, a cry of “oops!”, a desire for two short sticks). To evaluate the quality of causal evidence, children take knowledge, ability, and intention into account. For example, infants and preschoolers distinguish intentional actions from accidental ones, and this leads them to make different causal inferences (Carpenter et al. 1998; Meltzoff 1995). Preschoolers prefer to learn new causal relations from knowledgeable rather than ignorant causal agents (Kushnir et al. 2007). Children also treat causal evidence differently when a demonstrator is explicitly teaching them (Bonawitz et al. 2009; Rhodes et al., in press). This evidentiary link is not limited to passive observations – it influences and interacts with the evidence children generate themselves through play. Thus, when children get ambiguous evidence from another person, they privilege evidence from their own past actions (Kushnir et al. 2009), or are motivated to explore further to generate new evidence (Schulz & Bonawitz 2007).

Other research suggests that children break traditional domain boundaries when learning about people, as well. For example, infants use contingency detection (Shimizu & Johnson 2004) or violations of contiguity (Saxe et al. 2007; Spelke et al. 1995) to infer the presence of a psychological agent when other cues to agency are absent. Toddlers and preschoolers infer other people’s preferences based on violations of random sampling, not merely positive regard and enthusiasm (Kushnir et al. 2010). Children may use statistical cues to track other individual regularities, such as personality traits (Siever et al., under review). They also readily track social regularities, such as norms and group characteristics (Kalish 2002; Rhodes & Gelman 2008).

From her earliest social experiences, the child scientist is engaged in a dynamic process of hypotheses formation, evidence-gathering, and theory change. The adult moralist, on the other hand, is asked to reason about a single instance of human behavior. The adult must therefore rely on her existing knowledge – knowledge acquired through this early learning process. We now have a better sense of where this knowledge begins; recent studies show early understandings of empathy, fairness, help, harm, and a host of moral precursors (e.g., Hamlin et al. 2007). Knobe’s analysis encourages us not to stop with domain-specific characterizations of knowledge. Instead, we should broaden how we view evidence from human actions to include their moral and normative dimensions, and investigate how these early evidential links give rise to later moral asymmetries

in reasoning. This approach leads to interesting questions for research with adults, so long as we carefully distinguish between reasoning based on existing knowledge and the process of learning something new. When adults learn, for example, how do moral asymmetries change in response to further evidence? Is the evidence itself evaluated asymmetrically?

To conclude, while it may be wise at times to abandon the separation of disciplines, it seems premature to draw conclusions from Knobe's experimental data about the process by which they are integrated. To better understand this process, we need to look at learning at all ages, and continue research connecting moral development to both causal learning and social cognition.

Scientists and the folk have the same concepts

doi:10.1017/S0140525X10001809

Neil Levy

Florey Neuroscience Institutes, Carlton South 3053, Australia; and Oxford Centre for Neuroethics, Oxford, OX1 1PT, United Kingdom.

neil.levy@philosophy.ox.ac.uk

Abstract: If Knobe is right that ordinary judgments are normatively suffused, how do scientists free themselves from these influences? I suggest that because science is distributed and externalized, its claims can be manipulated in ways that allow normative influences to be hived off. This allows scientists to deploy concepts which are not normatively suffused. I suggest that there are good reasons to identify these normatively neutral concepts with the folk concepts.

Joshua Knobe has added considerably to our knowledge of the ways in which ordinary people attribute intentions and make judgments regarding causation. In this commentary, I do not want to criticize his claim that the competencies agents deploy in making these judgments are deeply suffused with normative influences. However, I will suggest that there are nevertheless grounds for regarding these competencies as distorting influences on our concepts. Our perfectly ordinary concept of causation (for instance), I suggest, is not normatively suffused. This is best brought out by thinking about science; I therefore begin with Knobe's claim that we ought not to understand folk judgments on the analogy of scientific hypothesis testing.

Knobe's claim that folk judgments are made in ways very unlike scientific hypothesis testing leaves us with a puzzle: Given that scientists are ordinary people too, how do *they* manage to engage in scientific research? If the relevant competencies are suffused with normative influences, how do scientists manage to free themselves of these influences (sufficiently well that they can identify them in the first place)? This question is important for several reasons, including, that if we can identify the means whereby scientists succeed in separating normative influences from the relevant judgments, we might all be in a position to make better *normative* judgments. At least on standard normative theories, our normative claims ought to follow *from*, rather than themselves cause, judgments of causation and intention; hence, separating out the normative from the non-normative might be a precondition of justified normative judgment.

So how do scientists manage to transcend the normative influences Knobe identifies? The answer is multifaceted, but an important part of it refers to the structure of the scientific enterprise. Science is an essentially distributed enterprise. The structure of a scientific community enables its members to compensate for the limitations and biases of individuals (Kitcher 1993). Individual biases can thereby be cancelled out; one scientist's bias toward a hypothesis will be cancelled out by another's against it. Of course, this cancellation process is powerless against the kind of normative influences Knobe identifies, as they are universal.

But the structure of science has a second property: it externalizes scientific knowledge. Since science, by virtue of its essentially distributed nature, requires that data and theories be available to a multiplicity of researchers, they must be presented in a format that makes this possible, and that requires externalization. Once theories and data are externalized in this way, they become available for manipulation using formal techniques, and these techniques are designed to be impervious to the normative influences Knobe identifies. They can also be manipulated through the use of methods such as double blinding, which can also serve to filter out normative influences.

One implication of the forgoing is that the finding that ordinary people are not best understood on the model of scientists is unsurprising: no one is a scientist alone. An agent can be a scientist only as part of a community of researchers engaged in systematic inquiry. The contrast between scientific judgments and folk judgments is therefore misplaced: The contrast is not between different modes of thinking so much as between different ways of manipulating mental representations; one individualistic and the other deeply social.

An important implication is that there are grounds for seeing the competencies that agents utilize in making judgments as distortions of their concepts. We do not wish to say that scientists are mistaken in making causal judgments that are not normatively suffused. We therefore should not see the concept of causation as constituted by the structure of the competencies Knobe has elegantly uncovered. Scientists are members of the folk, and their onboard competencies are identical to everyone else's, yet they understand their causal judgments, *qua* scientists, as deploying the ordinary concept of causation, not a theoretical innovation. I suspect that given the choice between the concept of causation used in science and one that is *explicitly* normative, ordinary people would also choose the former, providing further evidence that scientists use the ordinary concept.

In saying this, I take issue neither with Knobe's arguments in favor of the view that our competencies are themselves normatively suffused, nor with his correlative claim that the rival view (according to which moral judgments bias our application of our concepts) is false. I am accepting that normative influences figure into the relevant competencies, but I am claiming that nevertheless we need to distinguish between these competencies and the relevant concepts, even though we probably derive the concept from the competency (via some process of idealization). The concept of causation is normatively neutral, even though ordinary people deploy the concept using competencies that are normatively suffused.

It may be that we can dissociate the normatively neutral concept from the normatively suffused competencies only by externalizing and distributing our application of our concepts. We can hope to deploy our concepts better by becoming more like scientists. Doing so does not involve changing our onboard competencies – that may be a task that is beyond us – but instead requires that we alter the context in which we deploy them. By dividing and distributing cognitive labor, and by designing institutions that filter out the normative influences, we may become better reasoners, both in the normative and the non-normative realms.

Putting normativity in its proper place

doi:10.1017/S0140525X10001810

Tania Lombrozo and Kevin Uttich

Department of Psychology, University of California–Berkeley, Berkeley, CA 94720.

lombrozo@berkeley.edu

uttich@berkeley.edu

http://cognition.berkeley.edu/

Abstract: Knobe considers two explanations for the influence of moral considerations on “non-moral” cognitive systems: the “person as moralist” position, and the “person as [biased] scientist” position. We suggest that this dichotomy conflates questions at computational and algorithmic levels, and suggest that distinguishing the issues at these levels reveals a third, viable option, which we call the “rational scientist” position.

In this elegant and provocative article, Knobe summarizes a growing body of work suggesting that moral considerations influence a range of “non-moral” judgments, from mental state ascriptions to causal ratings. Knobe offers two interpretations for these data: (1) his preferred view of people as “moralists,” and (2) the traditional position of people as intuitive “scientists,” albeit poor ones subject to moral biases. We unpack these options using Marr’s levels of analysis, and suggest at least one viable alternative, which we call the “rational scientist” position.

In Knobe’s “person as moralist” position, “moral considerations actually figure in the *competencies* people use to make sense of human beings and their actions” (sect. 1, para. 7, emphasis added). In contrast, the “person as scientist” position claims that the “fundamental” capacities underlying these judgments are analogous to processes in scientific inquiry (sect. 2.2, para. 2). Both positions, as laid out by Knobe, involve a distinction between the “fundamental” or “primary” aspects of a cognitive system and those that are “secondary.” Knobe suggests that to account for the data, the scientist approach must claim that moral considerations play a secondary role, biasing judgments that are fundamentally scientific.

Examining these positions in terms of Marr’s levels of analysis (Marr 1982) reveals two different questions at play: one at the computational level, about the function of the cognitive system in question, and one at the algorithmic level, about the representations and processes that carry out that computation. For an advocate of the moralist position, the computational-level description of a cognitive system appeals to a “moralizing” function (perhaps evaluating people and their actions), and the algorithmic level is merely doing its job. For an advocate of the “biased” scientist position that Knobe considers, the computational-level description appeals to a scientific function (perhaps predicting and explaining people’s actions), but the algorithmic level is buggy, with moral considerations biasing judgments.

This leaves two additional options (see Table 1). First is the “biased moralist” position, with a “moralizing” function at the computational level, but a buggy algorithm. Without a fuller computational-level analysis that provides a normative account of the

judgments the algorithmic level *should* generate, this position is hard to distinguish from the “non-biased” moralist.

Second is the “rational scientist” position, which we advocate for some cognitive systems (Uttich & Lombrozo 2010). According to this position, a given cognitive system has a scientific function at the computational level, and the algorithm is just doing its job. To account for the slew of data Knobe cites, an advocate for this position must explain how moral considerations can influence judgments without threatening claims about the system’s function (at the computational level) or the efficacy of the processes that carry out that function (at the algorithmic level).

In a recent paper (Uttich & Lombrozo 2010), we attempt precisely this for ascriptions of intentional action. The cognitive system in question, broadly speaking, is theory of mind: the capacity to ascribe mental states to others. Traditionally, this capacity has been conceptualized as analogous to a scientific theory, with the function of predicting, explaining, and controlling behavior. At the computational level, this puts the traditional picture in the “scientific” camp. But what are the implications for the role of moral considerations in carrying out this function? Knobe seems to assume that moral considerations have no legitimate role in this picture. But we argue the reverse: that accurately inferring mental states can in fact require sensitivity to moral considerations, particularly whether a behavior conforms to or violates moral norms.

Here, in brief, is our argument. Norms – moral or conventional – provide reasons to act in accordance with those norms. For example, a norm to tip cab drivers provides a reason to do so. Observing someone conform to this norm is relatively uninformative: We can typically infer knowledge of the norm, but not necessarily a personal *desire* to provide additional payment. In contrast, norm-violating behavior can be quite informative, particularly when other mental-state information is lacking. If we believe a person knows the norm, then observing that person fail to tip a driver suggests an underlying preference, desire, or constraint that is strong enough to outweigh the reason to conform. This same logic applies to Knobe’s chairman vignettes (sect. 3.1). When the side effect of the chairman’s actions helps the environment, he is conforming to a norm, and the action is relatively uninformative about his underlying mental states. When he proceeds with a plan that causes environmental harm, the action is norm violating, and allows us to infer underlying mental states that support an ascription of intentional action.

Our aim here is not to elaborate and marshal evidence for this position; we direct interested readers to Uttich and Lombrozo (2010). Rather, we hope to populate the space of possible positions and call attention to what seem to be distinct computational- and algorithmic-level assumptions lurking in the background of Knobe’s target article. Knobe argues against various versions of the “biased scientist” position, but does not consider the “rational scientist” position. Like the two “moralist” positions, the biased and the rational scientist positions can be difficult to distinguish, and require a more fully specified computational-level description with a corresponding normative theory to identify which judgments stem from buggy versus non-buggy algorithms.

Knobe infuses normativity into folk considerations, painting a picture of people as moralists. But distinguishing the four positions we identify (Table 1) may actually require appeals to normativity in the generation and evaluation of empirically testable theoretical claims. In other words, we must appeal to normativity as theorists, regardless of whether or how we do so as folk. We suspect that Knobe avoids this framing as a side effect of other commitments and a preference for process-level theorizing. Whether or not it was intentional, we think it is a mistake to collapse computational and algorithmic questions. We hope future debate can restore normative questions to their proper place in scientific theorizing, whether the folk are ultimately judged scientists or moralists.

Table 1 (Lombrozo & Uttich). *Four possible positions to account for the data Knobe cites demonstrating an influence of moral considerations on non-moral judgments, such as mental state ascriptions and causal ratings. The positions are expressed in terms of Marr’s levels of analysis, with one of two computational level functions, and algorithms that generate the judgments they do either as a result of their computational level functions (non-buggy) or because they are biased by other (e.g., moral) considerations (buggy).*

Four positions to account for the data Knobe cites		Computational Level Function	
		Scientific	Moralizing
Algorithm	Buggy	Biased Scientist	Biased Moralist
	Non-buggy	Rational Scientist	Moralist