Free Will and the Folk: Responses to Commentators*

Shaun Nichols**

Experimental research on folk intuitions concerning free will is still in its infancy. So it is especially helpful to have such an excellent set of commentaries, and I greatly appreciate the work of the commentators in advancing the project. Because of space limitations, I can't respond to all of the comments. I will focus on just a few issues that emerge from the comments that I think are especially promising for illumination.

1. Conditional analyses

In my contribution, I maintain that people regard choice as indeterminist. This is partly based on the work that Joshua Knobe and I have done on lay intuitions concerning determinism. We presented subjects with descriptions of two universes, one of which is deterministic about everything, the other of which is determinist about everything except human decision making. We found that the vast majority of subjects responded that the latter universe is most like ours. Knobe and I interpret this as evidence that people in our culture have indeterminist intuitions about choice. Eddy Nahmias challenges our interpretation by pointing out that, given how we describe the universes, our interpretation depends on presupposing that the lay notion of "could have done otherwise" expresses something incompatible with determinism. However,

* Thanks to Joshua Knobe, Eddy Nahmias, Paulo Sousa, and Manuel Vargas for discussion of some of the issues in this reply.

** Department of Philosophy, University of Arizona. Email: snichols@email.arizona.edu

as Nahmias notes, the compatibilist tradition offers a different account of the expression "could have done otherwise." This tradition offers *conditional* analyses of that expression. The idea is that when someone says that an agent could have done otherwise, what this really means is that *if* the conditions (or laws) had been different, the agent *would have* done otherwise.

In light of various counterexamples, conditional analyses are widely rejected in the contemporary free will literature (for review, see Kane 1996, 52-8). I find the counterexamples compelling, but I've learned (e.g., from Knobe 2003) not to assume lightly that my own intuitions are representative of the folk. One might, however, explore the conditional analyses empirically. I recently conducted a pilot study on the topic.[1] 75 undergraduates at the University of Utah were given the following vignette:

> On 4/13/2005, Bill filled out his tax form. At precisely 10:30 AM, he decides to lie about his income. But of course he didn't have to make this decision. Bill could have decided to be honest.

The subjects were then asked to judge whether a sentence sounded right or wrong (on a scale from 3 to -3). One group got the following sentence, modeled on conditional analyses:

> Bill could have decided to be honest at 10:30, 4/13/2005, **but only if** some things had been different before the moment of his decision.

The other group got the following sentence, modeled on unconditional analyses:

> Bill could have decided to be honest at 10:30, 4/13/2005, **even if** nothing had been different before the moment of his decision.

My prediction was that people would tend to give higher ratings for the sentence modeled on unconditional analyses than for that modeled on conditional analyses. This prediction was borne out ($t(73) = -2.464$, $p < .05$, two tailed); subjects were more likely to judge that the unconditional sentences sounded right (mean score for unconditional sentence,

[1] The basic method used here is borrowed from Knobe (2004).

+.76; mean score for conditional sentence, -.58). While not decisive, of course, this helps to shore up the philosophical consensus that conditional analyses do not reflect the folk notion of *could have done otherwise*; as a result, this helps to sustain our claim that our results provide evidence people regard choice as indeterminist.

2. Folk indeterminism: children

Several commentators raise questions that bear on the child's concept of free will. Wierzbicka makes perhaps the most radical challenge, suggesting that the questions I've put to children concerning free will are "culturally inappropriate". To be sure, most children have not heard anything quite like the question "If everything in the world was the same right up until she chose to steal, did Mary have to choose to steal?" However, nearly all children answered the same way to this question, while they showed less consistent responses to the physical cases. I will immediately concede that there are worries about how to interpret these responses given the complexity of the questions. Still, the fact that children respond differently to different cases suggests that the questions tap into *some* difference in the way the child thinks about the cases.

Set aside the complex questions I asked kids in which they were supposed to hold the past fixed. Let's consider the simpler case in which children were asked whether the person or the thing *could have done something else*. The results on this experiment were utterly unambiguous. Every single child said that the person could have done something else and nearly every child rejected this option for the thing. This set of results is much clearer than the results with the more complex question, and it strongly suggests that children have a facility with the expression "could have done something else".

The ecological evidence, limited though it is, lends further support to this. In the CHILDES database,[2] we find examples of adults using

[2] The CHILDES (Children's Language Data Exchange System) database for child language research is a collection of transcripts from everyday verbal behavior of several children and their families. The database was initially established to study children's language, but it has been an invaluable resource for studying a number of features of child psychology (see, e.g., Bartsch & Wellman 1995).

"could have done otherwise" language with children. In one case, after the last yogurt is taken, the mother says "You could have left it for me" (MacWhinney & Snow 1990). In another case, the child takes a toy from one adult and gives it to another. The first adult says "You could've asked me for it" (Bloom 1970) And in yet another case, the child throws a toy airplane, and an adult says, "You didn't have to throw the airplane" (Bloom 1970). We also find children spontaneously using this expression. After bumping into his younger brother, a child says "You could've watched out Marky. You just should've moved" (MacWhinney & Snow 1990). So the available evidence indicates that children are exposed to the expression, "could have done otherwise", they spontaneously use it, and in the lab they respond systematically to questions about whether an agent could have done otherwise. Children (at least in our culture) have *some* concept of "could have done otherwise".

Let's assume that the child's concept of "could have done otherwise" is incompatibilist, as I've suggested is the case for the adult. If that's right, we're left with the perplexing question of acquisition. How does the child come to fluency with the idea that agents "could have done otherwise"? As Charles Kalish notes in his commentary, a key source of information here is explicit language. So, children get the notion of *could have done otherwise* partly as a result of hearing their parents say things like "Susan, you could have called me to say you would be late" or, as in the case above, "You could have left the yogurt for me". In his commentary, Kalish also observes that modal terms like "ought" get use in importantly different ways – sometimes in the context of agency and sometimes outside of that context. Kalish notes that the child has to figure out which interpretation to give to the term "ought" in context, and to do this, the child must use his understanding of psychological causation. This seems quite plausible, and I think that if we follow Kalish's lead and consider a broader range of modal language, we get a new and illuminating perspective on the acquisition of an indeterminist notion of choice.

Although I previously (2004a) suggested that the key folk distinction with respect to determinism is between choice events and physical events, I now think that this might mistake where the distinction really lies. For even apart from choice events, the child often hears and uses a notion of possibility that is naturally interpreted as in conflict with determin-

ism. Here again are some real examples of modal language use from CHILDES:

Father: You could fall and get hurt Ross.
Ross (4;2): No. Not if I hold on to here and here I won't.
Father: You could . . . It's dangerous (MacWhinney & Snow 1990).

Ross (2;7): Marky [a younger sibling] might fall (MacWhinney & Snow 1990).

Father: . . . I guess you can't play football then.
Abe (3;9): I can too I can too I can too.
Father: Would it help you if we got a helmet?
Abe: Yeah.
Father: How would it help?
Abe: It would hit the helmet instead of me.
Father: And what would happen then?
Abe: I'd not be hurt except it could still land on my nose (Kuczaj & Maratsos 1975).

Adam (4;2): Paul [a younger sibling], you might cut yourself on this (Brown 1973).

It's natural to interpret these uses of *possibility* as in conflict with determinism. That is, it's natural to read Ross as saying that it's a genuinely open possibility that Marky will fall (or not fall), and Adam as saying that it's a genuine possibility that Paul will cut himself (or not). There are various determinist-friendly ways to read the language of possibility, and perhaps one of these ways is the best interpretation of children's modal language. But it's worth noting that some of most familiar ways to tame modality in philosophy don't look at all promising. One way to preserve a notion of *possibility* while being neutral about determinism is to treat *possibility* as a deflationary kind of epistemic possibility, on which to say that *p* is possible is to say "For all I know, *p* will happen". But this seems an implausible interpretation. When Adam says that Paul might cut himself, it seems unprincipled to maintain that Adam really just means "As far as I know, Paul will cut himself." Similarly, when Ross' father says "You could fall" and then repeats, "You could . . . it's dangerous", it's doubtful that Ross would interpret his dad as merely reporting on epistemic possibility. The simple epistemic possibility

interpretation is even less plausible when we move to statements about past possibilities. Parents say things like "You could have broken the lamp!" And kids come to use language this way as well. In CHILDES we find Ross (at age 5) saying that he climbed on a shelf and "It could have fell on us." (MacWhinney & Snow 1990). Similarly, Abe (at age 3) is asked what would have happened if they had seen a bear when camping, and he replies "we could have been killed" (Kuczaj & Maratsos 1975). Obviously he doesn't mean "For all I know we were killed." Thus, the simple epistemic gloss fails to provide a general account of children's judgments of possibility. There are, of course, more sophisticated ways to construe the notion of possibility as consistent with determinism. For instance, in an influential paper on free will, JJC Smart defends a compatibilist reading of "could have done otherwise" by drawing an analogy with other uses of the expression, outside the context of agency (Smart 1961). Suppose you drop a plate which stays intact, and you observe, "it *could* have broken". In planting the seeds for his compatibilist agenda, Smart indicates that what we mean when we say "the plate could have broken" is the following: "within the ranges of possible initial conditions covered by possible cases of 'dropping,' the known dispositional characteristics of the plate do not allow us to rule out the proposition 'it will break'". This formulation would allow us to interpret the children's use of "possibility" in a compatibilist way, but it doesn't exactly strike one as the most natural interpretation of Ross and Abe.

Obviously it would be too great a task to undertake here to try to demonstrate that the child has an indeterminist notion of possibility even outside the context of choice. The point of the foregoing was simply to make this seem a plausible hypothesis. If the child has an indeterminist notion of *possibility* outside of choice contexts, we get a new avenue for exploring the acquisition of the concept of free will. For now it seems that the child's indeterminism about choice is part of a more general indeterminism about possibilities.

There are, of course, enormous difficulties in explaining the acquisition of folk concepts. In the case of folk modal concepts, it's especially tempting to think there's an innate contribution. (Indeed, modal concepts were notoriously difficult for classical empiricists to accommodate.) But in lieu of a real story about acquisition, we can at least consider a likely function of the child's concept of possibility. Elsewhere I've sug-

gested that a primary function of modal concepts is to represent risk and opportunity (Nichols forthcoming). This was based on an informal review of modal talk in the CHILDES database. It is a salient fact about everyday conversation that children and parents tend to use modal language largely to convey information about risks and opportunities. As parents, when we point out risks to our children, what matters to us is making sure that our children are safe, that they avoid unnecessary risks. We typically don't care about trying to communicate a carefully qualified notion of possibility that is neutral about determinism. Similarly, when we point out opportunities to our children, we aren't concerned to get them to hedge the modal notions in a compatibilist way. We want to stress the options before them. In short, when we alert our children to risks and opportunities, deterministic explanation is pretty much the last thing on our minds, or theirs. In this light, it should not be surprising if our notion of possibility fails to be nuanced in a compatibilist fashion. The concept of *possibility* can serve the key function of representing risks and opportunities quite well without any compatibilist subtlety.

3. Folk indeterminism: culture

Manuel Vargas and Anna Wierzbicka both suggest that it's difficult to know whether the results on folk indeterminism would extend to other cultures. I entirely agree, and I find this the most interesting avenue for future exploration. Wierzbicka makes the intriguing observation that for many languages, including Polish, there is no way to translate claims like "did Mary have to choose to steal?". Of course, this doesn't show that people in other cultures lack a nondeterminist concept of free will, nor does Wierzbicka suggest otherwise. Rather, she suggests that the project of exploring folk intuitions about free will exploit the "Natural Semantic Metalanguage" (NSM) approach that she pioneered. The NSM project seeks to uncover semantic primitives that are present in all languages. This is a fascinating project, and I should be very interested to learn whether NSM approaches would reveal cultural similarity or diversity on whether choice is conceived as indeterminist. As a point of optimism, I would note that CAN (construed as 'potentiality') is among the 61 primitives thus far collected (Goddard & Wierzbicka 2002, 73), and

a number of philosophers have maintained that "can" is one of the terms that expresses indeterminist free will. Here's Van Inwagen:

> In ordinary English, the concept of metaphysical freedom finds its primary expression in simple, common words and phrases, and not in the grand, abstract terms of philosophical art that one is apt to associate with metaphysics . . . One of the simple words that expresses the concept of metaphysical freedom in English is 'can'.

Interestingly, another of the 61 semantic primitives collected by the NSM project is MAYBE, construed as "possibility" (Goddard & Wierzbicka 2002, 73). Thus, the NSM project indicates that the key modal notions I invoked in section 2 are widely present across cultures. Obviously careful work would need to be conducted to discern whether the notions of CAN and MAYBE that NSM claims are semantic primitives carry the indeterminist implications, and I very much hope that this work will be done.

Vargas suggests that in addition to the possibility of cultural differences in intuitions about free will, it's also possible that religious upbringing has an important impact on intuitions about free will. This shows up the need for a rather different experimental endeavor. For the possibility that Vargas raises would mean that even *intra*-culturally, there might be systematic variation in intuitions about free will. Again this is an extremely interesting avenue. On the currently available theories of how we acquire the belief in free will (e.g. Wegner 2002, Greene & Cohen 2004, Nichols 2004a and above), religion plays no important role. Rather, the extant theories all invoke relatively simple psychological or developmental tendencies. So if religion is a powerful predictor of intuitions about free will, that will be a key bit of evidence to reshape our acquisition theories.

4. Incompatibilism, epiphenomenalism, and catastrophe

Vargas recounts a study showing that physiological explanations of behavior undermine responsibility attributions more than psychological or "experiential" explanations of behavior (Montessoro et al. forthcoming). Vargas' interpretation is that a physiological explanation provides one source of folk unease about responsibility. But Vargas would also main-

tain that the folk regard determinism as undermining responsibility (see, e.g. Vargas 2005).

Nahmias takes a more radical approach. He uses the fact that physiological explanations are undermining to provide an alternative explanation of the results that Joshua Knobe and I got on our studies of responsibility and determinism. When we asked the abstract question whether anyone could be "fully morally responsible" in a deterministic universe, subjects tended to say "no". Nahmias suggests that subjects might have given this answer not because they think that determinism conflicts with moral responsibility but rather because subjects interpreted our deterministic scenario as implying a kind of reductionism that renders psychological states epiphenomenal. Nahmias reports an instructive study he conducted in which subjects are presented with a scenario describing an alternate universe. For one group of subjects, the behavior of the inhabitants is given a neuro-chemical deterministic explanation; for the other group, the behavior is given a psychological deterministic explanation. Nahmias found that most subjects in the neuro-chemical condition denied that the agents deserved blame, while in the other condition, most subjects said that the agents did deserve blame.

As Nahmias notes, there are a number of ways that the incompatibilist might explain away the results. I want to offer one explanation for the difference between the results on his study and the results on the studies that Knobe and I conducted. One key difference between the studies is that Nahmias asks subjects whether agents in the universe "deserve credit or blame for their actions" whereas we ask whether agents can be "fully morally responsible". We chose our wording because incompatibilists (e.g. Galen Strawson) tend to acknowledge that there are notions of responsibility that *are* consistent with determinism. What incompabilists maintain is that the kind of strong moral responsibility that most people endorse is not compatible with determinism. Hence, to be as accommodating as possible to the incompatibilist, we used an expression that was deliberately strong. So one explanation for Nahmias' results in the psychological determinist scenario is that people are willing to grant that agents in a psychologically determinist universe deserve *some* blame and credit; but those same people might still maintain that the agents in that universe don't deserve the kind of *full* blame that is reserved for those presumed to have libertarian free will.

Now, what about the neuro-chemical determinist universe? Why do subjects tend to claim that those agents don't deserve blame or credit at all? Nahmias suggests that subjects interpret the neuro-chemical universe as one in which mental states are merely epiphenomenal.[3] The idea that our behavior is not caused by our mental states is truly, deeply disturbing. Leave aside issues about determinism and indeterminism, if our actions aren't caused by our mental states, then commonsense psychology is profoundly mistaken. We think that our actions are caused by what we intend, and our intentions are produced by our thoughts and wants. Epiphenomenalism trashes all of this. It is an instance of the catastrophe scenario envisaged by Jerry Fodor: "if commonsense psychology were to collapse, that would be, beyond comparison, the greatest intellectual catastrophe in the history of our species . . ." (1987, p. xii). Indeed, this is one place that nearly everyone – libertarians (e.g. O'Connor 1995), compatibilists (e.g. Smart 1961; Frankfurt 1971), hard determinists (e.g. Pereboom 2001), and revisionists (e.g., Vargas 2005) – can agree. If epiphenomenalism is true, our practices of attributing responsibility are irretrievably wrong. (That's why it is so fortunate that epiphenomenalism is false.)

5. Folk determinism, intuitive dualism, and kludges

Paul Bloom is one of the few developmental psychologists willing to claim that young children have philosophically ill-mannered commitments. In his fascinating book, *Descartes' Baby*, he argues that "we are dualists who have two ways of looking at the world: in terms of bodies and in terms of souls. A direct consequence of this dualism is the idea that bodies and souls are separate. And from this follow certain notions that we hold dear, including the concepts of self, identity, and life after death" (2004, 191). In his commentary for this issue, he suggests that the belief in indeterminist free will is another piece of this central dualistic worldview. He writes, "humans universally think of consciousness

[3] Of course, in that case, many philosophers would maintain that the subjects are wrong to draw the epiphenomenal conclusion.

as separate from the physical realm. Just about everyone believes, for instance, that when our bodies die, we will survive . . . And just about everyone believes in free will." I am simply delighted to be in such good company on the issue of whether people believe in indeterminist free will.

Although Bloom and I agree that people believe in indeterminist free will, he's skeptical about my claim that people (in our culture at least) also have determinist intuitions about decisions. Bloom raises several apt points against the experiment on determinism reported in my original article. Some of these objections, e.g., that the sample is small and homogeneous, can be addressed directly by conducting more experiments. A bigger concern, however, is the complexity of the instrument. As Bloom notes, the task is intellectually demanding and many subjects failed to abide by the conditions of the thought experiment. Bloom suggests that a simpler approach would be to explain the problem of free will and determinism to subjects and ask them what for their considered opinion. I think this would be a worthwhile investigation, but the drawback is that it might end up masking the natural intuitions that can be elicited from subjects in situations with limited disclosure. That is, the reflective, all-things-considered views people have about free will might obscure the natural intuitions that make the problem of free will so vexing. For my suspicion is that the problem of free will gets its firm grip on us precisely because it's easy to elicit intuitions on both sides of the issue. Different tasks will elicit different intuitions that cannot be reconciled. Of course in some cases, as is the case with judgments about determinism, the questions themselves have to be fairly complicated.[4] The best approach here, to my mind, would be to conduct structured interviews based on (for instance) the doppelganger thought experiment. One can then use various sorts of comprehension checks, as is common in developmental experiments. My prediction is that with appropriate comprehension checks, one would continue to get determinist responses on explanation tasks.

[4] Nahmias' commentary provides another approach to eliciting determinist intuitions. In his scenarios, the universe is recreated over and over. This likely primes people for prediction and explanation, and so would, I suggest, incline them more strongly to deterministic responses.

The idea that central philosophical problems are generated by conflicting intuitions is nicely (if perhaps inadvertently) illustrated in Derek Parfit's *Reasons and Persons*, which includes one of the most important treatments of personal identity in the 20th century. The chapter entitled "What We Believe Ourselves to Be" is immediately followed by a chapter entitled "How We Are Not What We Believe". The irony is that *both* chapters are entirely driven by intuitions. That is, Parfit attempts to refute our intuitive view of the self by recruiting more intuitions. This might be an acceptable strategy (after all, some of our intuitive commitments are more central to us than others), but the point is that we naturally hold conflicting intuitions that bear on personal identity. The conflicts are there, ready to be triggered.

The situation is similar in the case of free will. While it's intuitive that our choices aren't determined, it's also intuitive that a person's actual decision has to have an explanation. There has to be *some* reason why he made the choice that he did. There is a smattering of interview data that lends a bit of support to this. In his influential paper on the folk model of the mind, D'Andrade gives a summary report of interviews he conducted with Western subjects (D'Andrade 1987). When asked whether someone could do something for no reason at all, one subject answered, "Really no reason at all? I'd say there should be some reason somewhere." When asked if the reason could be trivial the subject responded, "Could be trivial, could be anything. But there should be a reason" (p. 133). D'Andrade also asked whether someone could fail to act on their desire in the absence of any competing desires: "If there wasn't a counterwish, could it be the case that he just didn't go even though he wanted to?" The subject responded: "That's like a contradiction. Because that doesn't make too much sense. There would have to be a reason why the person didn't do it if they wanted to do it. There'd have to be some reason. . . . It wouldn't be that they just wouldn't do it" (p. 132). It's natural to construe these answers – "there has to be a reason" – as reflecting deterministic presuppositions for explanation.

Obviously this is a very limited source of evidence. D'Andrade only interviewed only a handful of subjects, and there are sophisticated libertarian accounts that can at least partly accommodate the idea that every action has to have a reason (e.g. O'Connor 1995). But as I suggested in section 2 with respect to modal judgment, it's useful to con-

sider the function of the system we're investigating. If we reflect on the function of the explanation system, it begins to make sense that we would have an implicit commitment to deterministic explanations. Our explanation system is plausibly there because it helps us understand and predict the world (cf. Gopnik 2000).[5] An explanation system that is willing to give up, cry "Uncle", would be less effective than an explanation system that won't take "brute fact" for an answer. So when we turn our sights to the task of explaining a person's choice, we expect there to be some reason for any decision.[6]

Thus, I think there are good reasons to believe that people have determinist intuitions in explanatory contexts. But I suspect there's a deeper disagreement that underlies the divergence between Bloom and me on the topic of folk determinism. To see this, I want to draw on Stephen Stich's commentary. Stich contrasts the "elegant machine" view of moral psychology with the "kludge" view. On the first plank of the elegant machine view, our moral capacity "is subserved by an integrated set of rules or principles which . . . are designed to work smoothly together." According to the alternative view, that morality is a kludge, "morality is not subserved by a well integrated collection of rules but by a hodgepodge of psychological mechanisms that are often in competition with one another." Stich uses evidence from articles in this issue and elsewhere to argue for the kludge view of morality. Stich's kludge view strikes me as exactly right, and I think that it might be recruited to make a similar case against the elegance of intuitive dualism.

Return to the initial quote from *Descartes' Baby*, "we are dualists who have two ways of looking at the world: in terms of bodies and in terms

[5] Gopnik has a detailed story about the explanation drive. While I think much of what she says about the function of having an explanatory drive is quite plausible, I would not sign on for her view that that there is a domain-general phenomenology that characterizes explanation (2000, 311).

[6] Our thirst for explanation is illustrated in another philosophical context – the cosmological argument for the existence of God (crudely, the argument that God has to exist because otherwise there's no explanation for why the universe exists). The reason the cosmological argument has the pull that it does is largely because we assume that there *has* to be an explanation for why the universe is here. Philosophers enshrine the underlying intuition – that there has to be an explanation for everything – in the *Principle of Sufficient Reason*.

of souls. A direct consequence of this dualism is the idea that bodies and souls are separate. And from this follow certain notions that we hold dear, including the concepts of self, identity, and life after death" (2004, 191). Bloom's commentary suggests that we can add indeterminist free will to the list of consequences of dualism.

While I agree with Bloom that people are intuitive dualists, I have a different view of the direction of causation. Bloom indicates that the belief in immortality and free will are consequences of a prior commitment to dualism. I would maintain that the situation is reversed. We believe in dualism because we have intuitions of immortality and free will. That is, dualism is embraced because it fits with other aspects of our intuitive worldview, including the intuition that actions aren't determined and that we exist past biological death.

Let's consider three key strands of intuitive dualism: the natural inclination to believe in an afterlife (see especially Bering & Bjorklund 2004), the apparent separateness of the conscious and physical realms, and the intuition of indeterminist free will. Although each of these is important to intuitive dualism, there are intuitive considerations that pull against each strand. The salient fact of biological death raises the obvious possibility that we are *not* immortal, and the fear of death is a persistent theme throughout recorded history, stretching back to *Gilgamesh*. The known effects of drugs on consciousness (virtually a cultural universal) challenge the separateness of consciousness from the physical realm. And, as I've argued here, deterministic explanations compete with indeterminist intuitions about choice.

If there is all this intra-psychic conflict, why is dualism the clear cultural victor? Not, I suggest, because dualist theory is an elegant machine – a naturally developing set of integrated principles about the essence of the mind. Rather, dualist theory is a cultural achievement that takes succor from intuitions about immortality, separateness, and freedom.[7] Its cultural success might be largely explained by the fact that dualism, with its promise of an afterlife and a locus for moral responsibility, is motivationally attractive. That's the kind of view I *want* to believe. And there

[7] If that is right, then of course we do owe some other explanation for why we have intuitions of immortality, separateness, and free will.

is reason to think that views that are motivationally attractive enjoy a significant advantage in cultural evolution (e.g. Nichols 2004b).

If we ask why we have the moral system that we do, Stich maintains that the likely answer is that bits and pieces of our psychology get recruited and pasted together, drawing on quirks of cultural evolution. The same is true, I suggest, about intuitive dualism. Several naturally emerging intuitions get drawn together in the cultural manifestation of dualism. Like morality, intuitive dualism is a kludge.

REFERENCES

Bartsch, K, & Wellman, H.
1995 *Children Talk about the Mind*. Oxford University Press.

Bering, J. and Bjorklund, D.
2004 "The Natural Emergence of Afterlife Reasoning as a Developmental Regularity." *Developmental Psychology*, 40, 217-33.

Bloom, L.
1970 *Language Development: Form and Function in Emerging Grammars*. Cambridge, Mass: MIT Press.

Bloom, P.
2004 *Descartes' Baby*. New York: Basic Books.

Brown, R.
1973 *A First Language: The Early Stages*. Cambridge, MA: Harvard University Press.

D'Andrade, R.
1987 A Folk Model of the Mind. In D. Holland & N. Quinn (Eds.), *Cultural Models in Language and Thought*. Cambridge: Cambridge University Press, pp. 112-148.

Fodor, J.
1987 *Psychosemantics*. MIT Press.

Frankfurt, H.
1971 "Freedom of the Will and the Concept of a Person." *Journal of Philosophy* 68: 5-20.

Gopnik, A.
2000 Explanation as orgasm and the drive for causal understanding: The function, evolution, and phenomenology of the theory-formation system. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and Cognition*. Cambridge, MA: MIT Press, pp. 299–324.

Greene, J. & Cohen, J.
2004 "For the Law, Neuroscience Changes Nothing and Everything." *Philosophical Transactions of the Royal Society of London B*, 359, 1775-1785.

Kane, R.
1996 *The Significance of Free Will*. Oxford University Press.

Knobe, J.
2003 Intentional action and side effects in ordinary language. *Analysis*. 63, 190-193.

2004 Intention, Intentional Action and Moral Considerations. Analysis, 64, 181-187.

Kuczaj, S. and Maratsos, M.
1975 What children *can* say before they *will*. *Merrill-Palmer Quarterly*, 21, 87-111.

MacWhinney, B. and Snow, C.
1990 The child language data exchange system: An update. *Journal of Child Language*, 17, 457-472.

Monterosso, J., Royzman, E., and Schwartz, B.
forthcoming "Explaining Away Responsibility: Effects of Scientific Explanation on Perceived Culpability." *Ethics and Behavior* 15.

Nichols, S.
forthcoming Imaginative blocks and impossibility: An essay in modal psychology. In S. Nichols (ed.), *The Architecture of the Imagination*. Oxford University Press.
2004a "The Folk Psychology of Free Will: Fits and Starts." *Mind & Language*, 19, 473-502.
2004b Is religion what we want? Motivation and the cultural transmission of religious representations. *Journal of Cognition and Culture*, 4, 347-371.

O'Connor, T.
1995 Agent causation. In O'Connor, ed., *Agents, Causes, and Events: Essays on Indeterminism and Free Will*. New York: Oxford University Press. Pp. 173-200.

Pereboom, D.
2001 *Living without Free Will*. Cambridge University Press.

Smart, J.
1961 Free will, praise and blame. *Mind* 70, 291-306.

Vargas, M.
2005 "The Revisionist's Guide to Responsibility." *Philosophical Studies* 125: 399-429.

Wegner, D.
2002 *The Illusion of Conscious Will* (Cambridge: MIT Press).