

*We're All Folk: An Interview with Neil Levy
about Experimental Philosophy and Conceptual Analysis*

Neil LEVY and Yasuko KITANO^{*,†}

The following is a transcript of the interview I (Yasuko Kitano) conducted with Neil Levy (The Centre for Applied Philosophy and Public Ethics, CAPPE)¹ on the 23rd in July 2009, while he was in Tokyo to give a series of lectures on neuroethics at the University of Tokyo Center for Philosophy. The aim of this article is not to reconstruct the interview precisely. Following the referees' inscriptions, I edited his words for publication with his approval².

Yasuko Kitano (Y.K.) – I'd like to talk with you about experimental philosophy today. Recently, experimental philosophy, which is often called "x-phi" by younger practitioners, seems to be becoming better known in the Anglophone world. Kwame Anthony Appiah's article came out in *The NY Times* in December 2007, and David Edmonds and Nigel Warburton's appeared in *Prospect* in March 2009. And this June, a programme on this movement was broadcasted on BBC Radio 4.

According to Edmonds and Warburton (2009), there are three types of experimental philosophy. The first uses new brain-scanning technology and looks for patterns of neural activity when subjects are presented with philosophical problems. In the second type, philosophers conduct questionnaires to discover folk intuitions, e.g. Joshua Knobe's works. In the third, they carry out fieldwork; they observe how people behave in particular situations. For example, experiments as to how people behave after finding a dime in the return slot in a telephone booth.

Some philosophers exhibit doubt about the significance of the second type of experimental philosophy. Timothy Williamson, in an article contributed to *Proceedings of the Aristotelian Society*, says "If anything can be pursued in an armchair,

* The Institute for Sustainability Research and Education, Hosei University

† To contact the interviewer, please write to: yskitano@gmail.com

¹ He is no longer at CAPPE, but at the Florey Neuroscience Institutes. His broad research interests cover at least the following fields: free will and moral responsibility, philosophical psychology, moral cognition, practical epistemology, and neuroethics. He is editor-in-chief of the journal *Neuroethics*. For his major publications, see Levy (2002a; 2002b; 2002c; 2004; 2007; forthcoming 2011). For more detail, see: <http://neuro-ethics.org/>.

² I am grateful to the anonymous referees, Florian Cova (Jean Nicod Institute) and John O'Dea (University of Tokyo) for their helpful comments on an earlier draft.

philosophy can” and afterwards in his book, *The Philosophy of Philosophy*, says “In particular, on the account in this book, the method of conducting opinion polls among non-philosophers is not very much more likely to be the best way of answering philosophical questions than the method of conducting opinion polls among non-physicists is to be the best way of answering physical questions.” (Williamson 2007:7)

(Q1) Where do you think experimental philosophy’s significance is to be found, if anywhere?

Neil Levy (N.L.) – I think there is some. Well, the kind of experimental philosophy that I’m most interested in isn’t actually on that list. It’s a kind that rather simply tries to manipulate and observe behaviour. So for instance, the experiment that I described yesterday which I’m going to run is a cognitive load manipulation³. The idea is, rather than just to ask people what their views are, which give you their reflected views, to do things to bring out what their unreflective views are, which can be interesting because their unreflective views can be shown to influence behaviour. So particularly for ethics, I think this kind of experimental philosophy is very interesting, but I don’t think that very much of it is being done.

Well, I can give you one example of experimental philosophy that has been done in that mould. We don’t think of it as being experimental philosophy as it was not done by philosophers. But it’s about free will. So why not call it experimental philosophy. It’s the paper by Baumeister and Masicampo⁴. They gave their subjects a passage from Francis Crick to read. In the passage Crick is mocking of the idea of free will; he says “there’s no such thing as free will, and only a fool would believe that.” And then they gave their subjects and matched controls that hadn’t read that

³ In the handout of the lecture under the title of “Resisting Weakness of the Will,” Levy points out that there are two basic systems of cognitive processes: system 1 and system 2. System 1 consists of mechanisms that are fast, operate in parallel, and are undemanding of cognitive resources. System 2, on the other hand, consists of mechanisms that are slow, operate serially, and are demanding of cognitive resources. When the agent is under cognitive load (e.g. multitasking), system 1 processes pick up the burden and take over from system 2. The agent acts less intelligently (in more stereotypical ways) when her or his responses are generated by system 1. On the basis of this account of cognitive processes and other empirical evidence from cognitive and social psychology, he argues that weakness of the will is simply a special case of the depletion of system 2 resources, throwing the agent back on to system 1. What he is suggesting in the end is this: the behaviour we call weakness of the will exists but we ought to jettison the (folk psychological) concept “weakness of the will” in favour of a focus on the interrelationships between system 1 and system 2 because neither the explanatory purposes of psychology nor our practical purposes of increasing our ability to maintain self-control are well-served by retaining the concept.

⁴ Levy confirmed that he had referred to Vohs and Shooler (2008) and Baumeister et al. (2009). For the relating discussions, see Nahmias (2009).

passage a test to do on computer where they were trying to get the right answers. People were able to cheat on their test. They were told that there was a problem with the computer and that sometimes the right answers would appear on the screen and they'd get some warning before that happened, so that they could press the space bar when there was a warning, so they didn't get the right answer and so they could test fairly. And they found that the people who had read the passage by Francis Crick cheated much more. They allowed the answer to come up and then pretended that they had come up with it themselves. So they concluded that belief in free will is conducive to good behaviour. Now, I think that there are all sorts of problems with that experiment. But still it's clearly in the ballpark of experimental philosophy. It's not asking people about their views about free will. They weren't asked any questions while their behaviour was observed and manipulated. I think that's often more interesting.

But to come back to the survey style, I think there are some problems. I've got some sympathy for Williamson's view. I expect people to be confused about various things if they haven't thought about them and it's not particularly interesting to me to learn that most people are Libertarians, or most people are Compatibilists. When people are asked to justify this they justify it in terms of the burden of proof. And I'm very unconvinced that burden-of-proof arguments actually carry any weight whatsoever. I don't want to know who has the burden of proof. I want to know who has the better arguments. I don't believe we're in a situation where burdens-of-proof can be decisive. One side says we need to conclude that Libertarianism is true because Compatibilists haven't carried the burden of proof or the other way. That makes sense in a court of law but I'm not sure that the idea of a burden of proof makes much sense in philosophy. So, I have my doubts about the relevance in the free will case of this just asking people for their views. Eddy Nahmias's paper, the error theory one⁵, is interesting for the same reason why, say, Josh Knobe's research is interesting. Knobe finds that intuitions about causation in ordinary people are affected by the moral status of the action in question. I think that's interesting because the same process might be affecting philosophers' judgements about causation too. And in fact, I haven't read it yet, but somebody told me that Josh has written a paper with Christopher Hitchcock, a big name in causation, just coming out in *The Journal of Philosophy*. They argue just that philosophers are themselves affected by the Knobe effect⁶ in their discussion of cases. So there's something really interesting, in the same way as when a psychologist discovers reasoning bias, not because it tells us about the folk but because it tells us about everybody, because we're all folk.

And Nahmias's paper is interesting and provocative because what he's offering is

⁵ Nahmias and Murray (forthcoming).

⁶ See a series of his earlier papers: Knobe (2003a; 2003b; 2004; 2006a; 2006b).

an explanation of why philosophers are tempted to buy libertarianism. He suggests that maybe their initial worry isn't about determinism at all but about bypassing of psychological states, and then maybe they've argued themselves into thinking they're libertarians on irrelevant grounds. Now I don't know if that's true but it's certainly an interesting claim.

Y.K. – So, let's move on to the next question. Looking back at the history of Anglo-American philosophy in the past 100 years, conceptual analysis has been its dominant strain. “Philosophy of this kind considers,” Edmonds and Warburton say, “not so much how things are, but rather how we think about them: the way we carve up the world, the frontiers of meaning, of what makes sense.” Experimental philosophy is often thought as a movement challenging such armchair philosophy. But I think there is need to be cautious here. Knobe and Nichols say, in the introduction to the anthology *Experimental Philosophy* they edited, “As far as we know, no experimental philosopher has conducted an analysis of one concept in terms of another.” But some philosophers, e.g. Mark Balaguer⁷ and Anntti Kauppinen⁸, think that experimental philosophy is doing conceptual analysis on the basis of experiments.

(Q2) If that is true then can we not expect much impact from the second type of experimental philosophy?

N.L. – I don't think that experiments on the Knobe effect are doing conceptual analysis as well. They're not asking you about what our concept of causation is. They're trying to work out the psychological processes involved in judgements of causation. You can't blackbox the concept while you're doing that, nor do I think that the experimental philosophy on free will is conceptual analysis. I think it presupposes conceptual analysis, because you need some kind of analysis of the concept in order to be able to formulate your questions, but that's true in psychology too. I used to be more impressed by the difference between philosophy and the sciences until I started reading the sciences very seriously and discovered that they're in much the same boat as we are. If you read the papers in cognitive science, there's a lot of conceptual analysis going on. There's a lot of theory building going on. And a great deal of it doesn't look very different to what philosophers do. Philosophers have been doing what Jesse Prinz calls “empirical philosophy”⁹ for a long time. That's been the standard way of approaching philosophy of mind, for instance, or philosophy of biology. Conceptual analysis is constrained by seriously delving into the empirical data. That's exactly what psychologists and cognitive scientists do. They just do

⁷ Balaguer (2009).

⁸ Kauppinen (2007).

⁹ Prinz (2008).

more of gathering of data themselves, but you know, some of them don't do much gathering of data at all. Some of them are in the armchair just as much as we are. All Souls at Oxford have just appointed a psychologist to a position that involves no lab research at all. It's the first one to my knowledge, but it's a possibility in psychology. But you know, she's a psychologist. She's doing empirical psychology just in the way that people do empirical philosophy.

And, to get *ad hominem*, I don't think Balaguer is entitled to his claim, because what he wants to do, to look at human decision-making, requires just the kind of work to dissociate the processes that are part of judgement that the experimental philosophers are aiming to do. So I don't see how we would possibly begin to see what the nature of human decision-making is, except by looking at the ways in which various decision-making processes dissociate under experimental conditions. That's what people like Wegner, Libet, Gazzaniga and so forth do. So I don't see how it's somehow unscientific for philosophers to do it.

Y.K. – Now let's move on to Balaguer's paper¹⁰. Let me first refresh your memory of his argument. His point is that “the semantic component of the problem of free will” is, in a “nontrivial way” irrelevant to “the metaphysical component of the problem of free will”, that is, irrelevant to all substantive questions about human decision-making processes. He first argues that (i) the compatibilism question reduces to the “what-is-free-will” question, and then that (ii) the “do-we-have-free-will” question collapses into the “what-is-free-will” question and the “which-kinds-of-freedom-do-we-have” question (libertarian freedom, Humean freedom, Frankfurtian freedom, and so on). The assumption here is that the “what-is-free-will” question can be thought of as a question to which the answer is correct if it captures the ordinary-language (OL) meaning (or meanings) of the expression ‘free will’ (this is the OL View). He then argues that (iii) the “what-is-free-will” question is irrelevant to the “do-we-have-free-will” question irrespective of whether or not the OL view is correct.

(Q3) You introduced his paper in the blog the Garden of Forking Paths. Why did his argument draw your attention?

N.L. – So I'll answer that question and then talk about the claims it created. Why did I choose that paper? Well, he's a really interesting philosopher, in many ways. One of the reasons he's really interesting is most libertarians are religiously motivated – most of them are, I think, American Christians and most of them are motivated by a belief that free will is necessary to justify heaven and hell, and libertarian free will. And many of them are also, as Christians, believers that free will requires to

¹⁰ See note 7.

be, and is, in fact, actual because of non-physical properties. So many of them are agent-causationists. If you ever have the misfortune to go to a society of Christian philosophers meeting, you'll discover all these philosophers who are working on free will are agent-causationists and by the agent causal power they mean the soul or something like that. And some of them are good philosophers. Tim O'Connor has a view like that and he has some strange views but is a very good philosopher. There's no doubt about it. But to meet someone like Balaguer, who is a libertarian yet, as far as I know, completely unmotivated by religious concerns and tightly motivated by his reading of the sciences, is very unusual. So the fact is that he's a good philosopher but also that he's got such an unusual position. He's a libertarian of a kind, but he's highly motivated by his reading of neuroscience. I saw him having a debate with Adina Roskies. She's a neuroscientist originally. She has a PhD in neuroscience, and edited the journal 'Neuron' for a few years and then she went back and did a Philosophy PhD at MIT. She's recently published a paper with Nichols in *The Journal of Philosophy* called "Bringing moral responsibility down to earth" giving a 2-D semantics account of folk intuitions about free will¹¹. She's a very good philosopher as well as a very good neuroscientist. I saw her debate, a few years ago, with Balaguer arguing that neuroscience is consistent with indeterminism of brain processes. But that's part of the reason. He's so unusual. And the paper is a very provocative one in many ways. It's a challenge to the entire free will debate, and very relevant. So that's why I chose the paper. Now, when you go back to his claims, there were two things I wanted to say. (Checks the interviewer's original question to refresh memory).

Oh, yes, the first thing I wanted to say was about ordinary language. You can see why Balaguer thinks that experimental philosophy is continuous, because he takes conceptual analysis to be about capturing the ordinary language use of the term pretty well and experimental philosophers seem to be about capturing what the ordinary language use of the term 'free will' is. But I don't think that's quite right though. In fact I think that all the questions are about moral responsibility in the experimental literature and the reason is that I don't think that free will has a well-defined meaning in the folk vocabulary. Peter van Inwagen did some work few years ago—pretty sure it was van Inwagen—in which he looked at some big databases of classics of English Literature, Chaucer all the way through to 20th Century. And he looked for the occurrence of 'free will' and it was just about two in the entire corpus of English Literature. It just isn't a phrase in ordinary language. I suspect people will recognise it but not really attach a meaning to it other than as some kind of philosophical term. That's one thing. The other thing is, if you're trying to capture what you mean by 'free will', you don't want your use of 'free will' to diverge from folk concepts. I think that the way to go about that isn't to try and capture what

¹¹ Roskies and Nichols (2008).

the folk mean by 'free will', but instead to ask what concept makes sense of our free will practices. Now, what our free will practices are is controversial, because most philosophers would say they're about our responsibility practices, but some don't agree. Some think that free will is not about responsibility or not essentially about responsibility and van Inwagen's one of them. So, that makes it difficult but certainly you can tie your concepts to everyday meaning without trying to tie them to what the folk mean – I don't want to call it conceptual analysis though. So that's one of the things that I wanted to say.

The other thing is, and maybe this is as much relevant to your fourth question as your third. Can we in fact do the kind of work that Balaguer wants us to do? Can we really separate this semantic question from the metaphysical question the way he suggests? I don't think we can. I don't know what it would mean to study human decision making, to try and get at the facts about human decision making, except to know in a very theoretically saturated way. When we wanted to know what the processes are, we have to understand them in terms of our concepts and there's a relationship here between our concepts and the facts as we try to describe them because we may find that they don't fit into our pre-existing concepts very well. But we can only do it in a theoretically saturated way, and this is kind of a cliché in philosophy and science, but it's still true that our perception is theory-laden; so we can't do the description except in terms of our concepts.

And so what's metaphysically interesting here is, I take it, at least in part, is whether human decision making is such that we actually exercise control over our choices, we're actually the authors of our choices, we're the originators of our choices. To answer those questions, we can't just look at what's going on in the brain. We have to ask questions about those processes and whether they satisfy conceptual conditions that we lay down on those notions. The best example here is the luck objection. Balaguer is a libertarian, and he thinks he can avoid the luck objection. Alfred Mele has recently written a book about this¹². The basic idea is working with Robert Kane's event causal libertarianism. On Kane's view we sometimes face situations in which it's metaphysically undetermined whether we'll do the right thing or the wrong thing, and it's open to us to choose either way, and it's undetermined which way we'll choose. Mele says, what decides whether we'll choose one way or the other? Well, it can't be, by hypothesis, anything about our state of character, or our beliefs or our desires, or even how hard we're trying if it's all willings or volitions. It can't be any of that, because by hypothesis, holding all that constant, both of the options are metaphysically available to us. So, Mele says it can't be anything but luck which decides whether we choose one way or the other. The agent doesn't settle it, luck does. And he says you can't be responsible for choosing A rather than B, if

¹² Mele (2006).

it's just a matter of luck that you chose A.

Now there's a big debate here, Kane has various responses, and other people have other responses, and so forth. But Kane has given us a description of the brain. It's very speculative but let's assume that it's true. Now we know all the facts of the brain. We know, Kane says, that under certain conditions, when agents are torn between two options, there are competing neural networks which simultaneously attempt to initiate conflicting actions. This disrupts the thermodynamic equilibrium in the brain. I think it's laughable science, but anyway, let's assume it's true. This magnifies quantum mechanical level events to the level where they actually have an effect at the macroscopic scale of the brain and it becomes metaphysically open to us, which is, one way or the other, with the indeterministic event causal process settling which way we choose. That's a complete description of the brain, or at least it's complete enough. So now we know all the facts. We've got what Balaguer says is metaphysics, but we still don't know if we're free, because we still haven't got an answer to the luck objection, because you can't look in the brain and find an answer to the luck objection. Both of them are assuming an identical description of the facts, but they're still diverging on whether the agent is free. And I think the same is true for Balaguer's view. We can know all the facts, and still there's an argument to be had about whether the agent is free, and that argument can only be settled by – I don't want to say conceptual analysis – conceptual work. So I just don't think that we can do what he claims we can do with our theory, i.e. effective conceptual work, but even if we could, it would still lead us with a lot of further conceptual work to do. I think I might have moved on to question four.

Y.K. – In Balaguer's view, experimental philosophy addresses the “semantic component of the problem of free will” just as traditional conceptual analysis does; that is, it is irrelevant to discovering the nature of the world. But in my finding, surely we could use the results of subject polls to define the word free will, but what experimental philosophers are doing, at least in Knobe and his friends' brand of it, is figuring out folk intuitions and detecting the underlying psychological processes that generate them. You can say that this is about how things are, which is not a semantic question. So you could say this is about “how things are” and it might be needed to create a third category – not only “what is free will” question and “do we have free will” question, but the third category “what do we think of free will” question. By the term “think” I mean having certain psychological states.

N.L. – I think that's right, but I think there are no unique divisions. I do think that in so far as we conceptualise the world and we are conceptualising beings. We never stop doing semantics.

Y.K. – Analysing concepts is a bit different from analysing folk intuitions, isn't it?

N.L. – It is. But I think analysing folk intuitions is most interesting when it bears back on our own concepts. We start to wonder about if our own intuitions are being affected by the same processes we find in the folk. So, here's an instance. Whether it bears up I don't know, but Tamara Horowitz had a paper¹³ arguing that intuitions about double effects in ethics. The idea is that there are certain states of affairs it's wrong to intend to bring about, but you can bring them about as a foreseen consequence of an action, so long as you don't intend to bring them about. A classic example: it's wrong for a doctor to intend the patient's death, but a doctor can give the patient a lethal dose of morphine intending only to relieve pain, even though the doctor foresees that the patient will die. Horowitz argued that prospect theory, which is one of Kahneman and Tversky's many theories, explains our intuitions of double effect. According to prospect theory, in terms of how we value, we think that losses are worse than foregone gains. Since prospect theory is clearly irrational, you can frame the same problem in such a way as to generate conflicting options. They would show that double effect is irrational, at least. That was the argument. Frances Kamm replied¹⁴ and argued that they're not the same thing, but at least that's an instance where looking at the way in which folk conceptualise is of direct philosophical interest, because it might show that the way that philosophers have thought about a problem is itself the product of a psychological process that they can't acknowledge, that they would see as irrational.

Y.K. – So when we're conducting conceptual analysis, irrationality never comes up, is that what you mean?

N.L. – Well, typically we bracket all those kinds of questions, when we're doing conceptual analysis, we're trying to think about how we think about a certain part of logical space. One of the things we do is to look for conflicts in the ways we think about them and try to resolve those conflicts. So, I suppose what I'm calling for, the way we try to do conceptual analysis, is in part via reflective equilibrium. So we're engaged in a normal reflective equilibrium process of attempting to harmonise the various intuitions and dispositions that we have with regard to a particular part of logical space. What this gives us is the material for wide reflective equilibrium. Wide reflective equilibrium asks us to constrain our reflective equilibrium in terms of the findings of science, or relevant theories. I'm suggesting that what this gives us are some relevant theories in terms of which to constrain that conceptual analysis.

¹³ Horowitz (1998).

¹⁴ Kamm (1998).

What weight to give to those findings is an open question. But, in so far as we can take those results seriously, then they should be a genuine constraint on the conceptual analysis. That is, they have to be a fixed point. I think this is the real goal of experimental philosophy, which is to find evidence for a particular intuition to be unreliable. Because it gives us a fixed point in our wide reflective equilibrium.

Y.K. – So, what would you say about Stich’s criticism about intuitions? What is your attitude to that?

N.L. – I don’t know of how to do philosophy without intuitions. Nor do I think that anyone know how to do any intellectual enquiry without intuitions. So in a broader sense, all of conceptual work that includes work in physics and history and cognitive science and everywhere else, is reliant on intuitions, because intuitions about simplicity, about coherence between parts of a theory all play an indispensable role. This is a cliché again, but given a set of experimental evidence, there’s an indefinitely large set of theories which are open to us. We simply ignore most of that logical space, and then we have an argument about, in that small part of that logical space, which of the theories is better supported by the evidence, and what experiments would be crucial to settle between live options. And all of that we conduct – by “we” I mean all of us, scientists as well as philosophers – we conduct by reference to our intuitions about what supports what. So I don’t see us getting out of intuitions. Now I do believe that the more data the better, simply because the more data you have available the finer grained your distinctions can be, and this is the way of progress. The more items of data, the more precise the area of logical space you can delineate. But I don’t believe that we can ever dispense with intuitions.

Y.K. – Is what you’re saying compatible with the idea that our intuitions are open to criticism, open to empirical studies?

N.L. – Yes. Including something that psychologists often do, that is to compare responses of ordinary subjects in experimental situations which differ in something seemingly trivial. The dime study, for instance¹⁵. Something really trivial like finding a dime will affect people’s behaviour. I think that philosophers are doing an analogous thing when they engage in thought experiments. And, philosophy done well should do this. Suppose you have an intuition about a certain case, then first of all – this is a point that Frank Jackson has made – you could see publishing that as a kind of experiment. People like Stich criticise people like Jackson, and Stich has explicitly criticised Jackson saying he’s doing experiments within one subject, just himself. But

¹⁵ For one of the classical examples of this kind of study, see Isen and Levin (1972).

Jackson says “Think about the Gettier case, for instance, and you can see that as a kind of experiment.” Gettier has this intuition that there are cases of true belief that aren’t knowledge. But then he sends it off to a journal, and that’s a bit of an experiment – we’re testing now to see if the referees share the intuition. And then it’s published and now we’re testing to see at large whether philosophers have the intuition. And it turns out that it’s a really staggeringly successful experiment, that of all the philosophers who read the paper, as far as I know, only one has not had the intuition: Stephen Hetherington. Everybody else “Ah yes, he’s right.” So, you need to see knowledge in general, all of it, not just philosophy as a distributed process. That’s one of my common themes. That’s one thing. The second thing is that we do need to pay more attention to the minor variations than we have. One legitimate thing to do, and one important thing to do if you get a thought experiment that generates a certain intuition, is to try to see whether the features that generate that intuition really are the ones that are philosophically important. It might be something trivial, like the dime, that is doing the work. But I think, again, philosophy is a distributed process and we know a lot about, say, Gettier case, because people have come up with lots and lots of different variations, and they haven’t made the effect go away.

Y.K. – Fair enough. Thank you very much indeed.

N.L. – Thank you. It’s interesting material, isn’t it? It’s philosophical methodology, and experimental methodology, and thank you for giving me the opportunity to think about it because I think I’m getting clearer on it myself.

References

- Balaguer, M. (2009). The metaphysical irrelevance of the compatibilism debate (and, more generally, of conceptual analysis). *Southern Journal of Philosophy*, 47, 1–24.
- Baumeister, R. F., Masicampo, E. J., & DeWall, C. N. (2009). Prosocial benefits of feeling free: Disbelief in free will increases aggression and reduces helpfulness. *Personality and Social Psychology Bulletin*, 35, 260–268.
- Horowitz, T. (1998). Philosophical Intuitions and Psychological Theory. *Ethics*, 108, 367–385.
- Isen, A. M. & Levin, P. F. (1972). Effect of feeling good on helping: Cookies and kindness. *Journal of Personality and Social Psychology*, 21, 384–388.
- Kamm, F. M. (1998). Moral intuitions, cognitive psychology, and the harming-versus-not-aiding distinction. *Ethics*, 108, 463–488.
- Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations*, 10, 95–118.
- Knobe, J. (2003a). Intentional action and side effects in ordinary language. *Analysis*, 63, 190–194.

- Knobe, J. (2003b). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16, 309–325.
- Knobe, J. (2004). Intention, intentional action and moral considerations. *Analysis*, 64, 181–187.
- Knobe, J. (2006a). The concept of intentional action: A case study in the uses of folk psychology. *Philosophical Studies*, 130, 203–231.
- Knobe, J. (2006b). The folk concepts of intention and intentional action: A cross-cultural study. *Journal of Cognition and Culture*, 6, 113–132.
- Levy, N. (2002a). *Being Up-To-Date: Foucault, Sartre and Postmodernity*, New York: Peter Lang.
- Levy, N. (2002b). *Moral Relativism*, Oxford: Oneworld.
- Levy, N. (2002c). *Sartre*, Oxford: Oneworld.
- Levy, N. (2004). *What Makes us Moral? Crossing the Boundaries of Biology*, Oxford: Oneworld.
- Levy, N. (2007). *Neuroethics: Philosophical Challenges for the 21st Century*, Cambridge: Cambridge University Press.
- Levy, N. (forthcoming 2011). *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*, Oxford: Oxford University Press.
- Mele, A.R. (2006). *Free Will and Luck*, Oxford: Oxford University Press.
- Nahmias, E. (2009). Why ‘willusionism’ leads to ‘bad results’: Comments on Baumeister, Crescioni, and Alquist. *Neuroethics*.
- Prinz, J. J. (2008). Empirical philosophy and experimental philosophy. In: Knobe, J. & Nichols, S. (eds.), *Experimental Philosophy*, (pp.189–208). Oxford: Oxford University Press.
- Roskies, A. L. & Nichols, S. (2008). Bringing moral responsibility down to earth. *Journal of Philosophy*, 105, 371–388.
- Vohs, K. D. & Shooler, J. W. (2008). The value of believing in free will: Encouraging a belief in determinism increases cheating. *Psychological Science*, 19, 49–54.
- Williamson, T. (2007). *The Philosophy of Philosophy*, Malden MA: Blackwell.

(Received 2010.1.27; Revised 2010.6.7; Revised 2010.10.15; Accepted 2011.1.14)