

In: Christoph Lütge, Hannes Rusch und Matthias Uhl (Hrsg.): *Experimental Ethics. Toward an Empirical Moral Philosophy*. Palgrave Macmillan, New York 2014, 211–226

Experimental Philosophy is useful – but not in a specific way

Jacob Rosenthal, University of Konstanz, E-Mail: jacob.rosenthal@uni-konstanz.de

In this contribution I try to tentatively assess the relevance of experimental for systematic philosophy. Imagine that some philosophical claim is debated. Pros and cons are exchanged, thought experiments brought forth, implications are asserted, disputed, and evaluated, specifications and distinctions introduced etc. – and now the disputants receive information on how ordinary people assess certain scenarios associated with the claim. That is, they learn about “folk intuitions” on the topic. (In rare cases they may also learn about associated patterns of brain activity, but I will put this aside, as it raises different issues.) How should such findings affect the discussion? This is the question I am going to pursue here. The significance of experimental-philosophical results for philosophical debates could well depend on the debate in question, but as it happens, this is largely not the case – or so it seems to me. To my mind, a general diagnosis can be argued for.

1 Introduction: What is behind experimental philosophy?

In its methodological aspects, experimental philosophy is a branch of psychology. Thus, the question concerning its relevance for philosophy could be understood as simply being a case of asking if empirical findings of such-and-such kind are relevant to this-and-that philosophical claim. There is more to it, however.

First, as its designation indicates, experimental philosophy is a kind of philosophy. It could not be pursued largely independent of it, like other branches of empirical science. The surveys are connected to specific philosophical problems, and thus presuppose familiarity with theories and ongoing debates in philosophy on the side of the investigator. The scenarios and questions could not be framed, nor would the survey results mean much without intimate knowledge of these debates. Thus, the experimenters not only need philosophical along with psychological expertise, but there

is an asymmetry between these: The latter concerns the methodological aspects of the investigation, whereas the former supplies its aim and motivation.

Second, the test subjects are confronted with questions directly concerning (certain aspects of) the philosophical topic, framed in a comprehensible way. Surveys aim at finding out whether, e.g., ordinary people tend to a causal or rather to a descriptive theory of reference, whether their notion of knowledge is internalistic or externalistic, whether they tend to a deontological or to a consequentialist mode of moral reasoning, or whether they are compatibilists or libertarians concerning free will and determinism. It is not as if their behaviour in certain situations was observed in order to draw, e.g., conclusions for moral psychology potentially relevant to moral philosophy, but rather the subjects are simply *asked* certain questions that concern moral philosophy. This cannot be done in a straightforward way, to be sure, as people would not understand, but by confronting them with concrete scenarios described without technical vocabulary. It is as if, e.g., experimental economics did not investigate into the actions of subjects in experimental settings in order to test the empirical adequacy of behavioural models, but rather simply asked test subjects what *they* think how people behave under the relevant circumstances. The results could well be psychologically illuminating, but would have no bearing on the question of the adequacy of the economic models under consideration.

These peculiarities make experimental philosophy a truly unique enterprise and the question of its philosophical significance especially pressing. Its source is relatively plain, however. It is that in philosophical argument there is much (explicit or implicit) appeal to so-called *intuitions*. If one is not going to be content with implications, captured by conditionals like “if your concept of moral responsibility is such-and-such, then the person in this-or-that scenario is responsible for what she does”, at some point one has to flatly assert or deny that a certain case is a case of moral responsibility. The same goes for any other philosophical topic. The term “intuition” does not matter here – in this sense, its widespread use is a relatively recent phenomenon. The point is a systematic one: In order to make progress towards a definite stance on a philosophical question you need some anchoring in assertions that are not conditional in kind.

But how does one know that a certain carefully described case is or is not a case of knowledge, or reference, or morally required behavior, or free will? Well, it is just apparent that the respective predicate applies or does not apply to the case at hand. This, at least, would often be the claim. Such intuitions (whatever they are exactly – a bit

more on this later) provide a basis for philosophical reasoning. Consequently, they are not supposed to be mere subjective inclinations, but assumed to be *shared*, and not only by experts, still less by philosophers of a certain provenance (in which case the appeal to them in argument would be pointless), but rather by any reasonable person.

Thus, the role of intuitions in philosophy is sometimes compared to that of perceptions in the empirical sciences. They provide the data on which the theory is built, and if people did not agree on the data, the whole enterprise of empirical science or philosophy would collapse. I cannot investigate into this analogy further here, however. It is not easily assessable, for there are also clear differences: Unlike perceptions, philosophical intuitions may have abstract, general contents, they also arise in connection with purely hypothetical cases, and they probably cannot be calibrated in anything like the way our faculties of perception can. It may well turn out that there is no useful analogy after all.

Anyway, the intuitions in question are supposed to be common-sense- or everyday-intuitions. They concern “what we would say” about the case at hand. The tacit idea is that our concept of, e.g., reference or knowledge or moral obligation or free will commits us to certain judgments, so that everybody who masters the respective concept is bound to apply it to a clearly described scenario in the same way. There may be performance errors, to be sure, but by and large the conceptual competence should prevail (more on this later). Properly analyzed, the answers of the subjects reveal theoretical commitments implicit in our everyday ascriptions, and so contribute to the correct philosophical account of reference, knowledge, moral obligation, free will etc.

If this is the tacit assumption behind appeals to intuition in philosophical debates, it is evident that experimental philosophy is of great importance. Do ordinary people really make the assessments ascribed to them by philosophers who argue for a certain claim and back it up by appeals to intuition? How can we find out without actually conducting surveys? Psychological investigations are required. And they have to be done by people with philosophical expertise. It would be of little use, e.g., to put forward questions concerning the compatibility of free will and determinism to ordinary people if one has never thought about the different possibilities for explicating “determinism”, or if one simply takes it for granted that determinism precludes the ability to do otherwise than one actually does. If you do not know the respective philosophical debates, you are not sensitive to certain important distinctions and will most likely lump together different issues, so that you can neither formulate the questions properly, nor

assess the answers given by the test subjects. The point of experimental philosophy stands.

2 Appeal to intuition in philosophical argument and everyday intuitions

Before reviewing actual accomplishments of experimental philosophy, I elaborate on the fundamental matters a bit further. Three remarks are in order.

First, the sketched justification for experimental philosophy is tied to the idea that it is really common-sense-intuitions that are to be given argumentative weight in philosophical discourse, and thus, that the appeal to intuition (whether or not it is called by that name) is meant to refer to *these*. I should not like to claim that this is always the case, still less, that it must be so. What is in the background here is the view that philosophical issues concern – *inter alia* at least – the proper understanding of everyday concepts and their ordinary modes of application. It is not my business here to argue in favour of this understanding of philosophy; I just note that it is widespread. Along with it goes the tacit idea that philosophy cannot easily be revisionary. A philosopher like Plato, on the contrary, who distinguished “true happiness” or “true knowledge” from what the man in the street thought about happiness or knowledge, would have little use for surveys. In the Socratic dialogues everyday intuitions are raised all the same, but with the consequence or even the purpose of questioning them and dumbfounding the dialogue partner in order to make him reconsider. *This* kind of eliciting and dealing with ordinary people’s answers to philosophical questions is very different from what experimental philosophy does, and this difference reflects a partially diverging understanding of philosophy.

Second, the talk about intuitions is somewhat misleading, as it suggests a special source, mode or content of the judgments in question. There is a broad discussion about what philosophical intuitions, properly speaking, are, how much authority can be claimed for them, and how they are to be distinguished them from other mental states (see, e.g., the volumes DePaul and Ramsey 1998, Horvath and Grundmann 2012, or Booth and Rowbottom 2014). For reasons of space I do not enter this debate here, but note a dilemma. If “philosophical intuition” is given a distinctive sense, on the one hand, it is quite improbable that it covers all or most (explicit or implicit) appeals to intuition in philosophical discourse. The greater the extent to which it achieves that, on the other hand, the more open and vague it has to be. Whatever merits a clear-cut concept of

intuition may have, whatever epistemological function intuitions in some distinctive sense might fulfill, it is not intuitions in any such sense that are elicited by experimental philosophy. At the very least, there is no reason to think so. Presumably, no interesting kind of mental states is captured by the term “intuition” in its encompassing sense (see Williamson 2004 and 2007, ch. 7).

When laypersons are asked about certain scenarios, they probably answer quite spontaneously, and without doing much explicit theorizing, the background for which they lack anyway. This does not mean, however, that they do not make inferences, or conscious inferences – after all, they are supposed to make the case at hand clear to themselves, so there is plenty of time for episodes of reasoning. Moreover, nothing is known about a specific conceptual or modal content of people’s assessments, or the modal force people would be inclined to ascribe to their answers. We also do not know whether they provide mere guesses or firm convictions, whether or not their answers are based on persistent “intellectual seemings”, or anything of this kind. Thus, it is opinions in the broadest sense of the term the surveys may be said to deliver, and no more. Most importantly perhaps, there is no reason to think that what people say is pre-theoretical in any fundamental sense. There can be little doubt that the philosophical, religious, and scientific disputes of bygone ages and the various theories emerging from them have partly shaped what all of us, and philosophically untrained subjects in particular, think today.

So, the “folk intuitions” experimental philosophy surveys are nothing else than ordinary people’s assessments of certain cases, made after not too long a thought. If intuitions of this kind provide a basis for philosophical theory, it is a basis in the sense that we just happen to take certain things as preliminary and defeasible starting points. It is at best an open question whether the answers given by test subjects mirror intuitions in any epistemologically distinctive sense. This, however, does not diminish the relevance of the investigations. When an alleged common sense has a role to play in philosophical argument, it is simply a good idea to test whether the opinion in question is in fact widely shared by ordinary people. Whether you are prepared to call it “intuition” or not, experimental philosophy cashes out what is otherwise merely assumed. There are no fundamental objections against experimental philosophy that would not also count against the very appeal to intuition in philosophical debates – *if* only this appeal is indeed meant to refer to what “we” or “one” or “any reasonable person” would say.

It is understandable that several worries I can barely mention arise in connection with such surveys. The test subjects may misconstrue the issues at hand or misinterpret central notions. Their answers may depend on irrelevant features of the presentation or its context. Different scenarios may evoke inconsistent responses, in which case it is not clear which should be given priority. The answers people give might be influenced by doubtful pieces of theory or ideology. In short, people's responses may depend on all sorts of distorting influences that are bound to produce "performance errors". Thus, one may feel that they have to be debunked rather than given argumentative weight.

The answer to objections of this sort is the same as before. Even if it was granted that philosopher's judgments are considerably less subject to distorting influences of certain kinds, it would not matter. If you *really* distrust the "common sense" – and there may be plenty of reason to do so – then you may not appeal to common-sense-intuition or -usage in philosophical argument. Conversely, if you rely on it, you should respect what ordinary people in fact say when asked – unless they are somehow led astray. But again, this leading them astray must not be too easy. If there is such a thing as ordinary people's intuitions, they should be displayed quite consistently and robustly.

For the same reason, one should not be too quick in detecting distorting influences and performance errors. When it is really everyday intuitions and (implicit) theories that are at stake, the task is precisely to find out which features of a situation are considered to be relevant by ordinary people, and in what respect. This is achieved inter alia by varying scenarios. Claiming that people are misled by irrelevant features of the presentation or give inconsistent answers when the scenario is changed is often blatantly circular, simply assuming one's own assessments and distinctions in view of the matter.

To sum up, it is no use to claim that everyday intuitions are such-and-such, but that people, when actually asked, often answer differently. If the appeal to 'what we would say' is meant to carry argumentative weight, it has to introduce an independent element to the debate. Everyday intuitions may not be just hypothetical assessments, made by idealized reasoners or under counterfactual circumstances. There is no point in claiming that *if* people only thought long and hard enough about the issue at hand, they would share one's own opinions. To know the actual judgments of people, one has to conduct surveys. This is the only way to seriously check the validity of appeals to 'what we would say', in contrast to 'what we *should* say'.

Third. Not all, but many experimental-philosophical studies seek to find out about the assessments of laypersons by asking several students that have not (yet) taken a course in philosophy. The students volunteer for participation; they are not selected in a way that is liable to yield a representative sample of the student sub-population, let alone of the population as a whole. Moreover, the total numbers of test subjects are rather small: Often just a few dozens, at most a few hundreds. This constitutes a serious methodological objection to experimental philosophy. In the introductory book by Alexander (2012), this embarrassment displays itself in an omission pattern: There are no exact data reported in the main text; the reader is informed about the experimental parameters only in the endnotes. But not about the total numbers of participating subjects! This information is left out altogether. The reader can only roughly estimate those numbers from the reported proportions and significance levels.

Accordingly, most experimental-philosophical results have to be treated with caution, or rather, given credit. There is no reason to be particularly confident that the results would generalize when representative surveys were conducted. This problem has to be mentioned, but I am not going to dwell on it. Large and representative surveys are costly and laborious. What experimental philosophers do is not too far away from the general practice in psychology, the social sciences, or experimental economics. Even here, test subjects are very often simply taken from data bases of student volunteers, and in comparable numbers. For the sake of argument I assume that the results we deal with would generalize.

3 Actual achievements of experimental philosophy – a case study

The surveys done by experimental philosophy have so far yielded a variety of interesting phenomena. Just to recall some examples (which are all contained in the volume Knobe and Nichols 2008): The side effect effect (or Knobe effect) shows that opinions about the intentionality of some behaviour strongly depend upon the moral assessment of the agent and the context of the behaviour in question. Sensitivity to Gettier cases seems to vary with cultural context, so that it would be considerably more difficult to convince East Asian people that the traditional account of knowledge is defective than Westerners. Also, East Asians seem to favour a descriptive account of the reference of names over a causal one, contrary to what Westerners do. Several studies were conducted on the question of the compatibility of free will or moral responsibility and

determinism, with different scenarios and varying explications of determinism. These will serve as my prime example.

One of the most discussed philosophical questions concerning free will and moral responsibility is whether or not they are compatible with determinism (perhaps depending on the *kind* of determinism). The debate is very complex, of course, and several different positions are taken within each camp, but both sides draw to some extent on (alleged) everyday intuitions. Incompatibilists claim that free will and moral responsibility are tied to possible alternative courses of action, that these are precluded when everything that happens is fully determined in advance, and that compatibilists therefore have to dream up surrogate senses of “can do otherwise” that may look like as if they could do the trick for them but are ultimately pieces of sophistry. Compatibilists, on the other hand, feel that the question as to whether or not determinism is true is primarily one of natural science, that no matter what its results and theories are we would and should go on with our usual practices of praising, blaming and holding responsible, as well as uphold the corresponding judgments, and that it is muddle-headed to suppose that social practices and reactive attitudes towards one another should depend on what we think is going on on some fundamental level in the constitution of matter. This rough-and-ready sketch should do for our purposes. It is fair to say that both sides try to connect to common sense in some way.

Thus, it is a good idea to investigate into the intuitions of laypersons on the matter. In the surveys of Nahmias et al., varying majorities gave compatibilist answers. In several scenarios, the idea of determinism was spelled out in different ways, and actions of different kinds taking place under deterministic circumstances were depicted. Test subjects were asked whether the described agents were morally responsible for what they did and whether they acted of their own free will. Across the board, most subjects affirmed both of these, with majorities varying between around 60% and around 90%, depending on the case at hand (see Nahmias, Morris, Nadelhoffer and Turner 2005, 2006).

Other studies on the same matter had somewhat diverging results, however. Nichols and Knobe found that if the question of the compatibility of free will or moral responsibility and determinism is asked in an abstract way (with some characterization of determinism), most people answer negatively (around 85%). When specific cases are portrayed, there are up to about 70% compatibilist answers – but only with scenarios that are liable to provoke intense emotional responses, like a case of cold-blooded cruel

murder. When it is about cheating on one's taxes, most people give incompatibilist assessments (see Nichols 2004, Nichols 2006a, Nichols and Knobe 2007).

Only the last result mentioned is in direct conflict with what Nahmias et al. found out, but also the general tendency displayed in the studies of the two research groups undoubtedly points in different directions. (The discussion between the groups is explicitly conducted in Nichols 2006a, Nahmias 2006 and Nichols 2006b.) How could one explain the diverging results? First, different characterizations of determinism may come to mind. In all of the mentioned investigations, "determinism" and related terms were not explicitly used, because apparently laypersons either do not understand them or think that they rule out free will by definition (see Nahmias et al. 2006, section 3). While one could take this as evidence that the common sense favours incompatibilism, Nahmias et al. rather conclude that one has to circumscribe determinism. They have different ideas how to do this, but in all of them not only "determinism" is avoided, but also, to a conspicuously large extent, modal phrases like "is necessary", "has to happen", "inevitably" etc. Instead, they prefer to speak about predictability or causation or production or implications of initial conditions and laws of nature, whereas Nichols and Knobe make free use of phrases openly connected to necessity.

Second, there may be performance errors. According to Nichols and Knobe a majority of laypersons has incompatibilist intuitions, but these may give way to the desire to hold the agent responsible in emotionally charged cases. While this nicely fits the results of their own surveys, it does not explain why in the studies of Nahmias et al. people favoured the compatibilist answer even in neutral cases. Using the idea of performance errors in quite another way, Nahmias and Murray (2011) have argued that people give incompatibilist answers to certain scenarios because they confuse determinism with epiphenomenalism or fatalism, specifically, with the claim that what a person does is not causally or counterfactually dependent on episodes of practical reasoning, volitions, or intentions (see also Nahmias 2006).

Third, part of the difference between the findings may correspond to general differences in reply to abstract questions and concrete scenarios. That this is a phenomenon occurring in several domains was suspected and preliminarily investigated by Sinnott-Armstrong (2008), see also Knobe and Doris ms.. Many people are attracted to general principles regarding a certain concept that do not fit several of the standard applications of that concept, presumably because they are too demanding. General principles often set high standards that are partly ignored in practice, so that the

concept can be applied without much ado. One may think of ascriptions of intentionality, responsibility, or knowledge. While this may not be particularly surprising, I take it that it is a philosophically relevant topic not yet sufficiently explored.

Let's now put aside speculation about the sources of divergence between survey results and turn to an overall assessment of these. Although there is much more material to be reviewed I have not mentioned, it would only confirm the fact that there simply are no definite results as to whether or in which respects our everyday notions of free will and moral responsibility are compatibilist or incompatibilist. Still, the widely shared supposition is that they must be one of these, and that appearances to the contrary have to be explained away by reference to contextually caused distortions of various kinds (for an exception, see Feltz et al. 2009). Nichols (2006a) as well as Knobe and Doris (ms.) envisage varying criteria for free will and moral responsibility, but for different contexts. They still hold that, given a [specific](#) context, there is a shared folk intuition on the matter.

It is, however, this very supposition, that severe doubt is cast upon by the data. The most straightforward explanation of these is that there is no clear-cut everyday notion of free will or moral responsibility that allows an unambiguous answer to compatibility questions. The majorities vary greatly, and there are always considerable dissenting minorities. Diverging judgments may be traced back to performance errors, to different notions of free will and moral responsibility, or to vagueness in these notions. Or maybe the disagreement is not of a conceptual kind at all. It seems not unlikely that we face some hardly resolvable mix of these.

Even so, one can investigate further into these matters. Also, the very fact that there seems to be no unambiguous everyday belief on the compatibility of free will or moral responsibility and determinism may be surprising and was worth eliciting. That said, the philosophical discussion can largely go on as before, as the questioning of laypersons has not even remotely yielded anything like “what we would say”. The philosophical problems and options, and pros and cons are just the same and also as weighty as before. The only achievement of experimental-philosophical investigations concerning the systematic debate is that no one can claim to represent the “common sense”, thereby saddling his opponents with the burden of argument. This can be important indeed, and I will say a bit more about it later. But it is an unspecific benefit, restoring a symmetrical starting point. No position taken in the debate is particularly supported or undermined by the survey results.

4 Actual achievements of experimental philosophy – a tentative general assessment

The picture sketched in the last section generalizes. When one takes a closer look at any of the sub-enterprises of experimental philosophy, one has hardly the impression that the corresponding systematic debates within philosophy can or should be influenced by the empirical results to any considerable extent. The surveys yield nothing like a uniform or clear-cut picture. Very often, there is a majority's opinion, but the dissenting minorities are typically quite large. Thus, whatever question is asked, for each possible answer – in most cases there are just two alternatives – there are many people in favour of it.

This is somewhat concealed by phrases like “people (strongly) tend to” or “subjects are (much) more inclined to” or “are (very) likely to say this-and-that” which are often used to report findings of experimental philosophy. Officially, they are just shorthand for what is actually going on, but in fact they are not that innocent. They are liable to give the wrong impression that (almost) everybody tends to or is more inclined to or is likely to make a certain judgment. They suggest truths about each individual, and thus uniformity where none is to be found. Sometimes, qualifications are entirely dropped to yield statements like “In some contexts, people treat agency as indeterminist; in other contexts, they treat agency as determinist.” (Nichols 2006a, Abstract), which are misleading to the highest degree. It would be much better to avoid such phrases throughout and instead talk about what “many” or “more” or “most” test subjects said.

It is the quantitative distributions of survey results, not a problem one could detect from the outset, that deprives experimental philosophy of most of its potential influence on philosophical debates. You simply do not *learn* anything from those figures, as far as the systematic topic under discussion is concerned. Certainly a philosopher should not be bothered when her position is shared by “only” 20% or 30% of the laypersons, even if these numbers appear more robustly than in our paradigm case. She knows anyway that several trained and intelligent people disagree with her among her colleagues, but also, that many of them hold similar opinions. If she learns that the same is true for untrained subjects, what is she going to make of that? The question becomes all the more pressing if majorities are shaky, depending on wording, context, and

scenario, but this seems to be more of an additional problem. The core problem is that it is not at all evident that the numbers elicited by experimental philosophers tell a philosopher more than that there are several laypersons agreeing, but also several ones disagreeing with her.

A general, but unspecific benefit remains: The results of experimental philosophy remind us of something important that is all too easily forgotten in philosophical disputes, namely, that many philosophical questions, as well as most of the possible answers, have their roots in everyday thinking. They are not far-fetched inventions cooked up in response to hair-splitting problems the common person does not see. The plain truth is that ordinary people are as divided on the appropriate answers to these questions as are philosophers. The exact proportions do not matter much as far as the systematic problem is concerned, nor even on which side the majority falls. Only a large and robust majority – larger than 90%, say – would indicate a common sense, but suchlike seems to be lacking in all disputed philosophical topics.

Attempts to show, contrary to appearance, that there is a common sense, but that many people are misled by certain features of some of the scenarios at hand, are liable to make things even more complicated, like in our paradigm example. There is plenty of room for interpretation as to which applications of a concept are genuine and which constitute performance errors, as well as concerning the question whether different people apply different concepts (i.e., associate different concepts with the same verbal expression), or different folk theories tied to a shared concept, or concepts that are vague in crucial respects. It is far from clear that it is possible to unravel this, if only in principle.

A philosopher arguing for a certain position *already* has to confront the nagging problem why so many of his colleagues disagree with him. It may easily be as difficult to assess what exactly is going on when laypersons are divided on an issue. That was not clear from the start. There *could* have been a gain, i.e., a clear-cut picture of how laypersons think about the topic at hand might have emerged, but as a matter of fact, it has not happened. Asking laypersons thus proves to be just an epicycle of asking professional philosophers, with the additional disadvantage that it is not clear how many of them misunderstand the issue, are liable to make inconsistent judgments, what justification, if any, they would give for their opinions if they were criticized, how strongly they attach to them etc.

That it is the most straightforward reading of the results that people are simply divided on the matter is something experimental philosophers would not easily admit. Alexander, when discussing the “side effect effect”, writes: “How should we explain these minority responses? Typically, such responses are treated as noise resulting from some kind of performance error.” (Alexander 2012, p. 66) But why are they so treated? There certainly is no in-principle reason to consider the opinion of a minority of 20% or 30% to be “noise”. If so many of the test subjects can commit “performance errors”, then so can the majority of 70% or 80%. Are subjects misled by some features of the scenarios they confront? But everyday intuitions should display themselves robustly to deserve the name. As the experimenters did not deliberately try to lead their subjects astray, how could 20% or 30% of them simply miss the point? But *if* they could, why isn’t it the 70% or 80%?

Experimental philosophers are in general reluctant to admit that our everyday intuitions might be non-uniform or undecided, but they are happy to discover cultural or socio-economic or gender differences (for an overview, see again Alexander 2012, ch. 4). This is no surprise, since the former tends to deprive experimental philosophy of its systematic import, whereas the latter indicates that different cultures or socio-economic groups or genders favour different answers to philosophical questions or use crucially diverging concepts in the domain under discussion. This could undoubtedly trigger a lot of more studies. It is, however, differences of the very same kind that count as “performance errors” in the usual context that are supposed to reflect important conceptual or theoretical discrepancies in the latter. Kirk Ludwig, when discussing experimental-philosophical studies on criteria for knowledge ascription, remarks:

“It is not clear, of course, why the variation in responses across groups is to be counted as evidence for variation in norms while the variation within groups is not. If the variation within groups were treated as evidence for variation of norms, then of course the claim that there are different epistemic norms in different cultures or socioeconomic groups would have to be given up. We would get instead the conclusion that the same norms are present in the various groups, but adhered to by different proportions of people.” (Ludwig 2010, p. 161; see also Williamson 2007, p. 190–191)

It is of course possible to conclude from the fact that even after various investigations nothing like a clear-cut picture as to what “the folk opinion” on a certain topic might be emerges that further and deeper research has to be done. Perhaps questionnaires have to be supplemented by “neuroanatomical accounts of the cognitive

processes and mechanisms” and “evolutionary (or other teleological) accounts of the work that our folk concepts are supposed to be doing” (Alexander 2012, p. 69). Evidently, this opens a rich field for study, but it has nothing to do with experimental philosophy in its particular role. The case for it shifts ground if survey results are treated as psychological data among others. That results of empirical science can be relevant to philosophy is not under dispute. Information on cognitive processes, the evolutionary development of certain attitudes, concepts and practices, and on what “the folk” has to say about particular issues can be important in certain respects for certain areas of philosophy. There is no general diagnosis to be given about this. In this case, however, experimental philosophy amounts to nothing like a new or supplementary methodology for philosophy as a whole, but provides merely another kind of empirical data that may have to be taken into account by philosophers.

5 Conclusion

It seems that there are conflicting tendencies rooted within everyday thought as to what the precise application conditions of certain concepts are. What these tendencies exactly amount to is well worth investigating. Our concepts might be vague, or undecided in crucial respects, or people might associate different concepts with the same term, or the divergence might not be conceptual at all, but a matter of (implicit) theory. But whatever the right interpretation is, the idea to significantly enrich philosophical disputes by bringing in the common sense is ill-conceived. This “common sense” does not exist. Rather, the diverging tendencies within everyday thought explain to a good deal *why there are philosophical disputes at all*. Instead of helping to resolve them, the “everyday intuitions” constitute their background. Experimental philosophy helps to bring this to mind again, and herein lies its utility for philosophical investigations. There is, however, no confirmation or disconfirmation of a particular philosophical account of something by experimental philosophy.

Specifically important is only the following: In philosophical dispute, there is often the temptation to reduce or even get rid of the effort of bringing forth reasons that speak in favour of one’s own position. Instead, the opponents are saddled with the “burden of proof”, because their position is allegedly “counter-intuitive”. If this manoeuvre succeeds, the first party can sit back comfortably while the second has to work. The first party declares itself winner by default, so to speak (or, for that matter, in

case any serious doubts remain). I cannot assess the merits and drawbacks of such effortless fighting, but observe that at least one effort is required all the same: One has to argue that what the opponents have to say is *in fact* counter-intuitive in some serious sense, in comparison with one's own position.

The results of experimental philosophy accumulated so far show that such a claim should better not rely on a supposed common sense – very likely there is none. It is, on the contrary, quite probable that either party connects to some tendencies inherent in everyday thinking (and consequently also goes against other tendencies). Either position is intuitive in some respects or to some people, but unintuitive in other respects or to other people. Not that I think that undivided, clear-cut folk intuitions are never to be found – but it seems that there are none in philosophically disputed cases.

References

Joshua Alexander (2012): *Experimental Philosophy – An Introduction*. John Wiley & Sons.

Anthony Booth and Darrell Rowbottom (eds.) (2014): *Intuitions*. Oxford University Press.

Michael DePaul and William Ramsey (eds.) (1998): *Rethinking Intuition*. Rowman & Littlefield.

Adam Feltz, Edward Cokely and Thomas Nadelhoffer (2009): “Natural Compatibilism versus Natural Incompatibilism – Back to the Drawing Board”. In: *Mind and Language* 24, 1–23.

Joachim Horvath and Thomas Grundmann (eds.) (2012): *Experimental Philosophy and Its Critics*. Routledge.

Joshua Knobe and John Doris (ms.): “Strawsonian Variations – Folk Morality and the Search for a Unified Theory. In: John Doris et al. (eds.): *The Oxford Handbook of Moral Psychology*. Oxford University Press.

Joshua Knobe and Shaun Nichols (eds.) (2008): *Experimental Philosophy*. Oxford University Press.

Kirk Ludwig (2010): “Intuitions and Relativity”. In: *Philosophical Psychology* 23, 427–445. Also in: Horvath and Grundmann (eds.) (2012), 145–163.

Eddy Nahmias, Stephen Morris, Thomas Nadelhoffer and Jason Turner (2005): “Surveying Free Will: Folk Intuitions about Free Will and Moral Responsibility”. In: *Philosophical Psychology* 18, 561–584.

Eddy Nahmias (2006): “Folk Fears about Freedom and Responsibility – Determinism vs. Reductionism”. In: *Journal of Cognition and Culture* 6, 215–237.

Eddy Nahmias, Stephen Morris, Thomas Nadelhoffer and Jason Turner (2006): “Is Incompatibilism Intuitive?” In: *Philosophy and Phenomenological Research* 73, 28–53. Also in: Knobe and Nichols (eds.) (2008), 81–104.

Eddy Nahmias and Dylan Murray (2011): “Experimental Philosophy on Free Will – An Error Theory for Incompatibilist Intuitions”. In: Jesus Aguilar, Andrei Buckareff and Keith Frankish (eds.), *New Waves in Philosophy of Action*, New York: Palgrave-Macmillan, 189–216.

Shaun Nichols (2004): “The Folk Psychology of Free Will – Fits and Starts”. In: *Mind and Language* 19, 473–502.

Shaun Nichols (2006a): “Folk Intuitions on Free Will”. In: *Journal of Cognition and Culture* 6, 57–86.

Shaun Nichols (2006b): “Free Will and the Folk – Responses to Commentators”. In: *Journal of Cognition and Culture* 6, 305–320.

Shaun Nichols and Joshua Knobe (2007): “Moral Responsibility and Determinism – The Cognitive Science of Folk Intuitions”. In: *Nous* 41, 663–685. Also in: Knobe and Nichols (eds.) (2008), 105–127.

Walter Sinnott-Armstrong (2008): „Abstract + Concrete = Paradox“. In: Knobe and Nichols (eds.), 209–230.

Timothy Williamson (2004): “Philosophical ‘Intuitions’ and Scepticism about Judgment”. In: *Dialectica* 58, 109–155.

Timothy Williamson (2007): *The Philosophy of Philosophy*. Oxford: Blackwell.