

Science, Values, and Pragmatic Encroachment on Knowledge

Boaz Miller*

Forthcoming in the *European Journal for Philosophy of Science*

Abstract Philosophers have recently argued, against a prevailing orthodoxy, that standards of knowledge partly depend on a subject's interests; the more is at stake for the subject, the less she is in a position to know. This view, which is dubbed "Pragmatic Encroachment" has historical and conceptual connections to arguments in philosophy of science against the received model of science as value free. I bring the two debates together. I argue that Pragmatic Encroachment and the model of value-laden science reinforce each other. Drawing on Douglas' argument about the indispensability of value judgments in science, and psychological evidence about people's inability to objectively reason about what they care about, I introduce a novel argument for Pragmatic Encroachment.

Keywords Science and Values; Knowledge; Action; Subject-Sensitive Invariantism; Interest-Relative Invariantism; Intellectualism; Inductive Risk

1 Introduction

Mainstream western philosophy has usually kept knowledge and action in two independent domains. This separation has been recently challenged. According to a recent family of views in analytic epistemology, commonly dubbed "Pragmatic Encroachment" (PE), whether a subject has knowledge depends not only on the truth of her belief and its epistemic justification, but also on her social interests and pragmatic concerns (Hawthorne 2003; Stanley 2005; Fantl & McGrath 2009). In this paper, I regard PE as a substantive theory about knowledge, not merely about the attribution of knowledge in conversational contexts.

PE resonates with current arguments in philosophy of science against the received model of science as value free. Such arguments state that evidential thresholds for scientific theory acceptance inevitably involve social value judgments. This resemblance is not surprising as proponents of PE in analytic epistemology and current critics of value-free science in philosophy of science both acknowledge their debts to Rudner's (1953) classic paper "The Scientist *qua* a Scientist Makes Value Judgments" (Douglas 2009, 50-55; Fantl & McGrath 2010, 558). Yet, despite their common influence and similar views, these two debates hardly engage each other. This paper brings them together. I argue that PE and current criticism of value-free science mutually support each other. I devote the bulk of the paper to arguing that considerations against value-free science and empirical psychological evidence about the limits of human rationality provide a strong and novel support to PE.

In §2, I present PE and the parallels between PE and the current views on the role of social values in science. I argue that if PE is true, then the value-free ideal of science is false. In §3, I outline my main argument for PE, and distinguish it from existing arguments for PE. In the remaining sections, I unpack my argument. In §4, I review empirical evidence from experimental psychology and arguments in philosophy of science regarding the limits of rational inquiry and reasoning, which serve to support a key premise in my argument. In §5, I defend a moderate version of the "ought-implies-can" principle on which my argument relies. In §6, I argue that scientific theory

* B. Miller (✉)

Department of Philosophy and the Cohn Institute for the History and Philosophy of Science and Ideas, Tel Aviv University, Tel Aviv, Israel, 6997801. The Graduate Program in STS, Bar Ilan University, Ramat Gan, Israel, 5290002.

e-mail: boaz.miller@gmail.com

acceptance is not a merely pragmatic decision, and that accepted theories are also worthy of belief.

2 Pragmatic Encroachment on Knowledge and the Value-Free Ideal of Science

A standard assumption in epistemology is that knowledge is binary, i.e., either *S* knows that *p* or *S* does not know that *p*; but justification is a matter of degree, i.e., *S*'s belief that *p* can be more justified or less justified. Under this model, there is a threshold of *knowledge-level justification*, such that if the justification of *S*'s belief that *p* passes it, and *p* is true, and *S* is not gettiered, then *S* knows that *p*.

Let us call the traditional theory of knowledge “traditional intellectualism” (TI). What are the differences between TI and PE? The dispute between TI and PE is about what determines knowledge-level justification. TI is committed to two views. The first is *intellectualism*, which states that only traditionally conceived epistemic factors, such as the quality of *S*'s evidence for *p* or the reliability of her belief that *p*, are relevant to determining the threshold of knowledge-level justification. The second is *invariantism*, which states that the threshold of knowledge-level justification is always the same, i.e., it does not vary with context (Conee & Feldman 2004, 103-104, 296-297).

PE rejects intellectualism. It states that facts about *S*'s practical interests regarding *p* are also relevant to determining whether *S*'s belief that *p* passes the threshold of knowledge-level justification. Specifically, if *S* has high stakes regarding *p*, *S* is *ceteris paribus* in a worse position to know that *p* than if *S* has low or no stakes regarding *p* (e.g., Stanley 2005, Ch. 5). These need not necessarily be *S*'s *personal* stakes regarding *p*; scientists, for example, have responsibilities toward greater society (Douglas 2009, Ch. 4). PE accepts invariantism in the sense that knowledge-level justification is jointly determined by only two factors: the degree of epistemic justification of *S*'s belief that *p*, and facts about *S*'s practical interests regarding *p*. Namely, PE denies that knowledge-level justification varies with the context of knowledge attribution (see §3).

PE echoes the model of value-laden science, which has recently been extensively defended by Douglas. Douglas (2009, Ch. 5) distinguishes two roles of social values in epistemic evidential judgments: direct and indirect. The logical relations between knowledge, justification, and interests according to PE are analogous to the relations that Douglas draws between an epistemic judgment, evidence, and social values in their indirect-role capacity, respectively. In their *direct* role, values serve as reasons for making an epistemic judgment, such as theory acceptance. For Douglas, the *direct* role is *not* legitimate in the context of justification, as it amounts to wishful thinking. But in their *indirect* role, values determine the threshold level that evidence must meet for making a justified epistemic judgment by determining the levels of inductive risks we are willing to tolerate. Douglas identifies two types of inductive risks: wrongly accepting a hypothesis when it is false, and wrongly rejecting it when it is true. There is an inherent trade-off between them. The more we expose ourselves to one type of risk, the less we expose ourselves to another type. She argues that the indirect role is legitimate and required, because social values determine acceptable risks in a given context, and different social circumstances may legitimately require different balances between types of errors. When we value the possible consequences of a risk as mild, we lower the threshold level of evidence required for making a justified judgment, and when the consequences are acute we raise it.

There are historical connections between TI and the value-free ideal. When modern epistemology and philosophy of science emerged as independent fields in the

early twentieth century, the lines separating between the fields were blurry. The analysis of knowledge was a central project common to both fields. As Schlick (1917/1985, 4-5), the founder of the Vienna Circle, writes:

[T]he meaning of the word 'knowledge' seems so obvious that there is no need for a for a more detailed, careful elucidation [...] But circumstances may arise in which a more exact definition and elucidation of the word 'know' becomes absolutely necessary, where many who suppose themselves quite clear about its meaning would be altogether in error... [T]he theory of knowledge must first determine once and for all just what specific process the term 'knowledge' is to designate.

TI was taken as the starting point of this elucidation of knowledge. The thought was that the notion of justification, which occurs in the analysis of knowledge, could be best analyzed in terms of a scientific logic of confirmation:

[T]here is no doubt that in the sciences we really possess both knowledge and advances in knowledge. This implies that the sciences have at their disposal a sure criterion for deciding when genuine knowledge is at hand, and in what it consists (Schlick 1917/1985, 5; cf. Reach 1946).

This view, however, encountered difficulties. Within philosophy of science, logical empiricists failed to develop a defensible logic of confirmation, and with the introduction of the problems of theory ladenness of observation and incommensurability in the fifties and sixties, the debate took a step back, as a need arose to defend the very view that science indeed produces the epitome of rational knowledge. Gettier (1963) made clear that mere justified true belief is not knowledge. Consequently, the discourse about knowledge within analytic epistemology gravitated toward figuring out the bridging elements between mere justified true belief and knowledge. This turn in analytic epistemology was accompanied with a methodological shift from examining the features of our best scientific theories to probing into philosophers' intuitions about thought experiments in which subjects have or lack knowledge. The methodological divide between analytic epistemology and philosophy of science has remained to this day. Nevertheless, Schlick's point about the connection between the methods of science and the nature of knowledge is still valid. Specifically, considerations from the practice of science lend strong support to PE. And, as I will now stress, the debate about PE also has major implications to the debate about the value-free ideal.

The dispute between TI and PE is about the *types of conditions* that need to obtain for "*S* knows that *p*" to be true. TI holds that only conditions relating to the truth of *p* and *S*'s epistemic standing with respect to *p* need to obtain, while PE states that conditions concerning *S*'s practical standing with respect to *p* also need to obtain. If PE is right about knowledge in general, it is specifically right about *scientific* knowledge. Namely, if PE is right, scientific hypotheses must *inter alia* satisfy conditions concerning subjects' practical standing with respect to them if they are to rise to knowledge. The value-free ideal is incompatible with this outcome. The value-free ideal is a normative guiding epistemic conception of science that concerns the types of factors (only epistemic or both epistemic and practical) that ought in principle to influence scientists in their evidential judgments, including certification of scientific hypotheses as knowledge (Douglas 2009, 1). Simply put, PE states that knowledge is value laden,

which means that scientific knowledge is also value laden. If PE is right, the value-free ideal is false.

The incompatibility of PE and the value-free ideal clarifies the terms of the debate about the value-free ideal. We may distinguish between the types of conditions that scientific theories need to satisfy in principle for being certified as knowledge, and the types of conditions that current science actually satisfies. Once we make this distinction, and see that both PE and the value-free ideal concern the former, we can see why Betz's (2013) recent defense of the value-free ideal misses its target. Even if it is right, it does not save value-freedom as an ideal for science. Betz's argument is twofold. *First*, "there is a class of scientific statements which can be considered—for all practical purposes—as established beyond reasonable doubt" (218; emphasis added). *Second*, for the remaining class of statements that cannot be considered conclusive for all practical purposes, scientists need not allow values to indirectly influence their acceptance or rejection, but rather they can merely report them along with their error probabilities.

In §4.2, I will argue against the second claim. Note, however, that even if we accept it, Betz's argument does not vindicate the value-free ideal, but denounces it. The value-free ideal is *not* an empirical claim about the types of factors whose influence happens to be salient in the current state of science. At best, Betz's argument shows that a large class of scientific hypotheses trivially satisfies the practical conditions for their justified acceptance; namely, that the risks associated with their justified acceptance in all conceivable contexts are negligible. It does not follow that practical conditions are in principle irrelevant for making justified scientific epistemic judgments. On the contrary, Betz's repeated caveat "for all practical purposes" denounces the values-free ideal, because it explicitly states that practical concerns *are* relevant to justified theory acceptance. Once the caveat "for all practical purposes" is made, it remains merely an empirical contingent matter which hypotheses can be justifiably accepted for all practical purposes, most practical purposes, or just some practical purposes. Betz's argument misses its target; and the value-free ideal remains false, at least if PE is true.

In this section I presented PE and the value-free ideal. I stressed that if PE is correct about knowledge in general, it is specifically correct about scientific knowledge, which entails that the value-free ideal is false. This point has not been sufficiently acknowledged in the current debate about the value-free ideal.¹ In the next section, I outline my argument for PE, according to which the opposite is true as well: If the value-free ideal is false (due to Douglas' argument), then PE is right.

3 An Outline of the Argument for PE

Before I outline my argument, let me briefly review existing arguments for PE in order to contrast them with mine. There are two dominant lines of argument for PE. One line of argument for PE is from the semantics of knowledge attribution. Consider this example: Jack is in a train station in Ottawa, about to board a train to Toronto. He doesn't know if he is on the right platform. He asks a random normally-looking person whether this is the train to Toronto. The person says that it is (and it is). Does Jack know based on this testimony that this is the train to Toronto? According PE, the answer may depend on Jack's interests. If it is very important for Jack to arrive in Toronto on time, e.g., to attend an important meeting, then he does not know this, and needs to make

¹ This is although Douglas (2009) uses the word "knowledge" throughout her (2009) book, and declares at its outset that "we have no better way of producing knowledge about the natural world than doing science" (2009, 1). In addition, in her argument for the indirect role (2009, Ch. 5), she draws on Heil (1983), who – like the current debaters of PE – discusses knowledge and belief in general, rather than science.

further inquiries, such as asking a train-station employee. But if it doesn't matter to Jack so much, this testimony is enough to grant him knowledge. Either way, a belief that this is the right platform is equally justified. The relevant difference between the cases is Jack's stakes regarding this proposition.

Proponents of PE argue that competent speakers tend to attribute knowledge to subjects in low-stakes cases, but refrain from that in high-stakes cases. This includes self-attributions. When the stakes change, e.g., if for whatever reason it becomes crucial for indifferent Jack to arrive in Toronto on time, speakers recant their previous knowledge attribution. PE proponents argue that PE is the best explanation of this linguistic phenomenon; particularly, it is a better explanation than attributor contextualism, which states that in knowledge attributions, the threshold of knowledge-level justification changes with features of the *attributor's* conversational context, rather than with the subject's stakes (e.g., Cohen 1999; DeRose 1995). Critics of this line of argument question whether there is a genuine, stable empirical linguistic phenomenon that requires an explanation at all, and insofar as there is, whether PE is its best explanation (Pritchard 2006; Brown 2008; May et al. 2010).

It is important to note that proponents and many opponents of PE regard empirical evidence about knowledge ascriptions by competent speakers to normal subjects in various circumstances and conversational contexts as highly relevant for deciding the dispute between them. Such ascriptions refer to *ordinary standards of justification that are used in real-life contexts*, rather than standards of epistemic excellence. PE is a normative account of knowledge in that it attempts to analyze knowledge as an existing human normative institution, rather than the appropriate epistemic conduct of ideal or exceptionally able epistemic subjects.

Another line of argument for PE is from the knowledge-action principle, which has three versions: (KA1) knowledge is necessary for rational action, i.e., if one can rationally act on one's belief that *p*, one knows that *p*; (KA2) knowledge is sufficient for rational action; (KA3) knowledge is necessary and sufficient for rational action. If we assume fallibilism,² and KA1, and that a low-stakes subject may rationally act on her belief that *p*, but not a high-stakes subject, it follows that the low-stakes subject knows that *p*, while the high-stakes subject may not know that *p*. If we respectively assume KA2, it follows that the high-stakes subject does not know, while the low stakes subject may know (Stanley 2007, 200-201). The knowledge action principle has also been disputed, as well as what follows from it (Neta 2007; Brown 2008).

I do not wish to dwell on these arguments and their criticism, only note that my argument for PE is different. It is based on the nature of inquiry, the limits of human rationality, and the normative role of justification in epistemology. Thus, it does not fall prey to existing criticisms of PE. It also draws attention to fundamental questions underling the view, which have not been adequately addressed yet.

My argument for PE goes as follows:

- (1) In some cases and to some extent, scientists cannot avoid being influenced by pragmatic factors when making non-trivial epistemic judgments in the process of inquiry that ultimately leads to the justified acceptance of scientific theories.

² Fallibilism is, roughly, the view that a subject can know that *p* while acknowledging a remote possibility that *p* is false (Fantl & McGrath 2009, Ch. 1).

- (2) At least in many of the cases referred to in (1), the scientists who are influenced by pragmatic factors when making non-trivial epistemic judgments inevitably make non-trivial value judgments.
- (3) Therefore (from 1,2): In some cases and to some extent, scientists cannot avoid making non-trivial value judgments in the process of inquiry that ultimately leads to the justified acceptance of scientific theories.
- (4) If non-trivial value judgments are made in the process of inquiry that leads to the justified acceptance of scientific theories, then these theories embed and reflect those non-trivial value judgments, in the sense that these theories might not have been justifiably accepted had they been considered in light of alternative non-trivial value judgments, such as alternative weighing of competing values.
- (5) Therefore (from 3,4): In some cases and to some extent, justifiably accepted scientific theories embed and reflect some non-trivial value judgments, in the sense that these theories might not have been justifiably accepted had they been considered in light of alternative non-trivial value judgments.
- (6) Justifiably accepted scientific theories satisfy the justification condition for knowledge.
- (7) Justifiably accepted scientific theories satisfy the belief condition for knowledge.
- (8) The justification condition and the belief condition for knowledge are jointly sufficient for knowledge.³
- (9) Therefore (from 5-8), in some cases and to some extent, (scientific) knowledge embeds and reflects some non-trivial value-judgments, in the sense that it might not have been the same had it been considered in light of alternative non-trivial value-judgments.⁴

This argument requires unpacking and defense, which will be done throughout this paper. I focus on scientific knowledge, but my argument holds for any knowledge that is the product of methodic complex inquiry that involves non-trivial inferences, such as the knowledge produced by a police detective or an investigative reporter. In the next section, I defend the first premise of the argument.

4 Why the Influence of Interests and Social Values is Unavoidable

4.1. Douglas on the Limits on Inquiry

What is the basis for premise (1) of my argument; namely, that in some cases and to some extent, epistemic agents cannot avoid being influenced by pragmatic factors? This claim rests on two pillars. The first is Douglas' argument about the indirect role of values in science presented in §2. The challenge to value-free science from the indirect role of values was already posed by Rudner (1953), who argued that scientists cannot avoid weighing social risks against each other when deciding to accept or reject hypotheses. A standard reply to Rudner was that the need to evaluate research consequences arises only at the last stage of inquiry in the context of application; thus, scientists, who work in the context of justification, need not be those who evaluate

³ For the sake of this argument in this paper, I leave out of the discussion the truth condition for knowledge. For a discussion of the sense in which scientific theories may satisfy it, see Miller (2013, 1295-1296)

⁴ I thank an anonymous reviewer for improving my original reconstruction of the argument.

social risks (Jeffery 1966; McMullin 1983). This reply accords with the received model of value-free science. For example, Hempel (1965) agrees that social values play an indirect role, but maintains a strict division of labor: Science assigns various hypotheses probabilities, and only society assigns them utilities and decides which hypotheses to accept.

A novel feature of Douglas' work is showing that this reply is inadequate, because a strict separation between basic and applied science, or the contexts of justification and application, is unsustainable. Douglas argues that values penetrate deep into the context of justification, and affect various stages of research, such as study design, data analysis, evidence characterization, and evidence interpretation. Douglas (2000) supports her argument *inter alia* by a case study involving scientific research of the carcinogenic effects of dioxin, which illustrates the inevitability of making value judgments in various stages of inquiry prior to theory acceptance. She discusses a series of studies in which rats were exposed to dioxin and slides with their liver tissues were taken to determine if they had developed cancer. As part of the study, researchers needed to characterize the tissue samples. As it turns out, characterizing some slides, which is done by identifying certain visual patterns in them, is a subtle matter. Three different studies that used *the same slides as data* characterized some of them differently.

Douglas argues that this case illustrates the necessity and inevitability of the indirect role. Values should not influence the characterization of clear evidence. Clear cases of diseased tissues should be characterized as such and clear cases of healthy tissue should be characterized as such. But values ought to influence the characterization of borderline evidence. All other things being equal, in a society mostly concerned with the dangers of cancer, borderline slide cases should be characterized as diseased, and in a society mostly concerned with the economic burden of overregulation, they should be characterized as healthy. Such a practice reflects the types and levels of inductive risk that society is willing to take (Douglas 2000; 2009, 124).

The bottom line is that when research outcomes finally reach the context of application, they are already saturated with social value judgments, and reflect the various trade-offs between values that were made in the process of inquiry leading to them. Moreover, the indirect influence of social values is not a necessary evil that must be tolerated. Rather, it is necessary for achieving scientific objectivity because only it offers a non-arbitrary, principled, and relevant way to decide the various dilemmas that arise during inquiry and influence its outcomes.

In this paper, I take Douglas' reply to the standard objection to Rudner and subsequent developments of in the spirit of her account (Wilholt 2009; Intemann & De Melo-Martín 2010; Elliot 2011) to be robust, and do not offer new arguments for them. In the next subsection, I draw attention to another cause of scientists' inability to avoid the influence of values on their evidential reasoning, which is the psychological phenomenon of motivated reasoning. Motivated reasoning gives independent support to PE, and consequently also militates against the value-free ideal.

4.2. Psychological Evidence on the Limits of Human Rationality

The second pillar of my claim that in some cases and to some extent, epistemic agents cannot avoid being influenced by pragmatic factors is empirical psychological evidence about people's difficulty to objectively reason about what they care about. The influence of values on evidential reasoning is well studied in experimental psychology. The term "motivated reasoning" describes any process of reasoning that is affected by a person's

preference or desire concerning the *outcome* of the reasoning process (Kunda 1990, 480). Motivated reasoning is a species of confirmation bias (Nickerson 1998; Klayman 1995).

Motivated reasoning affects evidence assessment. People assess the same evidence differently based on their directional goals. Here are some representative examples. When coffee lovers read a scientific article claiming that caffeine was hazardous, they were less convinced by it than non-caffeine consumers. Sports fans were told that a previously winning team had lost a game. Fans of the team tended to see this as a mere fluke, while fans of the opposing team tended to see this as a turning point (Kunda 1990, 488-490). Proponents and opponents of capital punishment received the same mixed evidence about the effectiveness of that practice. Both regarded the evidence as reaffirming their prior beliefs. When presented with the same studies, scientists tended to deem as more methodologically sound the studies that supported their previous beliefs (Klayman 1995, 394-395).

These findings nicely accord with PE. They show that when the stakes are high, e.g., when coffee lovers encounter evidence suggesting they should give up their precious hobby, they raise the evidential thresholds required for knowledge, but they do not do so when the stakes are low or null. Motivated reasoning also casts serious doubts about Betz's suggestion, which was discussed in §2, that scientists report hypotheses they deem uncertain along with their error probabilities. The doubts are both about scientists' ability to distinguish between certain and uncertain hypotheses, and their ability to objectively determine the error probabilities. For the sake of the argument in this paper, I take these empirical findings to be robust.

Several objections may be raised against my taking these empirical findings to support a normative case for PE. Let me address them in turn. First, I take the psychological evidence as illustrating human irrationality, but this interpretation may be challenged. Feldman (2003, 162-164) is aware of the potential problems that empirical psychological findings that show human irrationality may pose to TI. Anticipating the line of argument I advance in this paper, he argues that psychological studies that reveal that people use incorrect rules of logic or probability do not necessarily show human irrationality, as people may follow different rules than the ones assumed by the researchers, which may still be rational although they lead to false beliefs in exceptional cases.

However, the cases to which Feldman refers, known as "cold cognitive biases", are not the ones on which my argument rests. I reviewed cases of "hot cognitive biases", where the source of irrationality is emotion, rather than misguided inference rules. Indeed, Feldman (2003, 165-166) admits that such cases cannot be rationally described, but argues that the existence of a limited range of cases in which people form irrational beliefs does not endanger normative epistemology or TI. These studies, however, are more alarming than Feldman depicts them, because they show that agents *systematically* fail epistemically; namely, in many and perhaps all cases where agents are motivated toward a certain outcome and the evidence allows for interpretive flexibility, they *will* likely reach a biased belief. This shows that TI may hold agents to higher standards than they can meet. (I discuss this point further in §5.)

It might be objected that these findings do not show that subjects *cannot* form justified beliefs in high-stakes cases, but only that some percentage of people *does not* form such beliefs. But there are three relevant respects in which the empirical evidence demonstrates an inability, rather than the mere prevalence of a specific form of reasoning. First, the process of motivated belief formation is, for the most part,

unconscious; namely, subjects are unaware of it. Normally, subjects do not realize that the belief-forming process is biased by their goals, and that they would probably form different beliefs in the presence of different directional goals (Kunda 1990, 482-483). Second, motivated reasoning is *involuntary*. Subjects do not willfully form motivated beliefs. In fact, even when subjects are trained to avoid certain biases, in a wide class of cases they nevertheless end up forming biased beliefs against their will. The success of debiasing techniques in experimental settings is modest at best, and there are good reasons to think that they are even less effective in real-life conditions outside the laboratory (Lilienfeld et al. 2009, 393-395). Third, it is hypothesized that a *psychological mechanism* is responsible for motivated belief formation (Kunda 2009). If this hypothesis is correct, then these findings reveal fundamental facts about the limitations of human reasoning. They, in a sense, identify the “normal working parameters” of rational human evidential reasoning. When these parameters are exceeded, namely, when subjects’ stakes are too high, people’s reasoning processes are no longer reliably rational.

It may be objected that even if these findings show that subjects cannot form justified beliefs in certain conditions, this failure does not have normative significance. To illustrate this claim, Conee & Feldman give the following counterexample:

A paranoid man might believe *without any supporting evidence* that he is being spied on. This belief might be a result of an uncontrollable desire to be a recipient of special attention. In such a case the belief is clearly epistemically unjustified even if the belief is involuntary and the person cannot alter the process leading to it (2004, 85; emphasis added).

Let us grant for the sake of the argument that the fact that the paranoid cannot help but form the belief that he is being spied on carries no normative force. Motivated reasoning differs from the paranoid’s case in three ways, which make it normatively significant.

First, the paranoid’s belief (according to Conee & Feldman’s own construction of the counterexample) is not responsive to evidence at all. The paranoid forms his beliefs regardless of his having no evidence for them. By their very nature, then, his beliefs are epistemically defective. By contrast, motivated subjects’ beliefs *are* responsive to evidence. The influence of directional goals on subjects’ evidence assessments is constrained. Motivated subjects are *not* at liberty to conclude whatever they want. Rather, empirical research finds that they are constrained by their ability to rationalize their reasoning. Subjects attempt to construct a justification of their desired conclusion that would persuade a dispassionate observer. They draw the desired conclusion only insofar as they feel they can do that. Subjects insist on appearing as following rational reasoning processes (Kunda 1990, 482-483). Thus, unlike the paranoid’s beliefs, motivated subjects’ beliefs are not categorically irrational or unjustified, but rather less rational or justified than non-motivated subjects’ beliefs. As Goldman (1999, 237), who makes a similar point, puts this: “Motivation driven reasoning has a robust reality, but there are limits to the epistemic damage it can inflict”.

The paranoid fails by forming beliefs regardless of the evidence, whereas motivated subjects allow directional goals to affect their evidence assessment, but they are still responsive to the evidence.⁵ Inasmuch as this is a failure, it is a failure of

⁵ If we reconstruct the example such that the paranoid suffers from a pathological influence on his evidential reasoning rather than his beliefs directly, the difference between the paranoid and normal people turns from a difference in kind to a difference in degree, and the example loses its normative force. It is obvious that because the paranoid is extremely influenced by motivating factors, his belief

justification, rather than belief. We therefore cannot draw lesson from the paranoid's case to the motivated subject case. As opposed to the paranoid, there remains a possibility that in some circumstances, motivated subjects' beliefs may still be sufficiently justified to reach knowledge-level justification.

Second, the paranoid suffers from an abnormal mental disorder, which affects his rational cognitive performance. When we deem his belief unjustified, the implicit standard we use is what a normal person whose cognitive capabilities are not damaged would believe in his place. By contrast, motivated reasoning occurs in normal people. It is *not* pathological. When we deem a high-stakes person's belief that *p* unjustified, the implicit standard we use is either a high-stakes subject, whose cognitive capabilities are higher than normal, or a subject whose stakes regarding *p* are low or none. Both these standards are inadequate. Proponents of TI cannot use the low-stakes subject as setting the epistemic standards to the high-stakes subject without begging the question; and epistemic standards for knowledge should not exceed normal human limits, thus the subject with higher-than-normal cognitive capabilities cannot be used for setting the standard. Hence, the objection from the paranoid man fails.

Against the last claim, proponents of TI may deny that epistemic standards should not exceed normal subjects' doxastic limits. I address this objection in the next section, where I defend the sixth premise of my argument.

5 Justified Acceptance of Theories Satisfies the Justification Condition for Knowledge

In this section, I defend the sixth premise of my argument, which states that justified acceptance of scientific theories satisfies the justification condition for knowledge. My claim is both empirical and normative. Namely, I claim that justification standards for knowledge do not and should not exceed the standard of justified acceptance of well-established scientific theories. This claim should not be controversial. Many proponents and opponents of PE already seem to accept it. This is because, scientists are regularly attributed knowledge with respect to their justifiably accepted theories in ordinary linguistic contexts, and as I noted in §2, most debaters of PE find knowledge attributions in ordinary conversational contexts as highly relevant empirical evidence for deciding the debate. If anything, the common view among philosophers and non-philosophers is that scientific knowledge is in *better* epistemic standings than non-scientific knowledge; as Longino (2002, 124) characterizes this position, "scientific knowledge is like ordinary knowledge except better".

However, because I am committed to neither the significance of knowledge attributions as evidence for deciding the case about PE, nor the superiority of scientific knowledge to other knowledge, I would like to defend this claim against the major objections that have been raised against it.

The rationale for the view that justificatory standards sufficient for knowledge do not and should not exceed justificatory standards for acceptance of well-established scientific theories is that epistemic standards for knowledge should not exceed normal human limits. This rationale is captured by the "ought-implies-can" principle, which is commonly regarded a truism.⁶ As the example of the paranoid in the last section

is not justified. But it does not follow that a normal person who is *moderately* influenced by motivating factors is also not justified. Drinking heavily before driving is irresponsible, but this does not mean that any amount of drinking before driving is also irresponsible. I thank an anonymous reviewer for pressing me on this point.

⁶ Cf. a central principle of Jewish law "we do not impose on the community a hardship which the majority cannot endure" (Babylonian Talmud, Baba Bathra 60b).

demonstrates, and as some have argued (e.g., Stocker 1971; Mizrahi 2012), however, this rationale has *prima facie* counterexamples and exceptions, particularly in the epistemic context. Yet for the sake of my argument, I need not defend a strong version of the epistemic “ought-implies-can” principle, according to which standards of responsible epistemic conduct must under no circumstances exceed a subject’s capabilities (hereinafter, *the extreme principle*). Rather, I need only defend a modest principle, according to which epistemic standards for knowledge should *generally* not exceed what people are ordinarily capable of achieving, and if there are exceptions, such as the paranoid’s example, they are relatively rare or unrepresentative (hereinafter, *the modest principle*). To be clear, the modest principle does *not* imply that there is nothing that is generally out of reach of our capacity for forming justified beliefs. There are cases in which what a subject who has done everything she can for reaching justified belief should do is suspending judgment or disbelieving.⁷

The modest principle is enough for counting most well established scientific theories as knowledge, because in the process of their acceptance, scientists for the most part do not exceed ordinary human capabilities. Despite a common myth, scientific knowledge is not produced only by geniuses, and scientists, *qua* individuals, are no more rational than the rest of humanity. “It should be obvious that the objectivity and the rationality of progress in science is not due to the personal objectivity and rationality of the scientist” (Popper 1981, 95).⁸ If, as I will argue in the next section, such theories inherently embed and reflect value judgments, then knowledge embeds and reflect the same value judgments, hence PE is true.

What possible objections are there to the modest principle? According to one objection, standards of justified belief are standards of excellence. According to a second objection, epistemic oughts that set justificatory standards for belief are role oughts, which circumvent the epistemic “ought-implies-can” principle. Both objections are raised by Conee & Feldman. I will address them in turn.

According to the first objection, epistemic oughts that are used in the analysis of justified belief set *standards of excellence*. According to this interpretation, proposed by Conee & Feldman (2004, 87), standards of justified belief are analogous to standards for an “A” grade, which may be too high for most and maybe all students. As such, it is fine that normal people cannot meet them. This interpretation is consistent with Kitcher’s (2011, 125) view that it is wrong to dismiss an epistemic ideal merely because it is difficult to implement, since without an ideal, we do not know where we want to go. On this interpretation, epistemic oughts do not imply can.

This objection, however, does not dispense with the epistemic “ought-implies-can” principle, but only relaxes it. While standards of excellent performance are higher than what most people can achieve, they must still be defined vis-à-vis what people in general can and are expected to do. If we follow Conee & Feldman’s analogy, an undergraduate student should not be expected to write a publishable research piece to get an “A” for his term paper. Even as excellence standards, epistemic oughts are *correlated* with people’s abilities. Indeed, Kitcher (2011, 125-130) agrees that we must at least have some initial ideas about how to start putting an epistemic ideal into practice for it to be a valid ideal. Thus, this objection dispenses only with the extreme principle, but not with the modest principle.

⁷ See Miller & Record (2013, 122-127) for further discussion of this point.

⁸ The view that scientific rationality (inasmuch as there is one) is not the individual rationality of scientists is largely uncontroversial in contemporary history, philosophy, and social studies of science. See Solomon (2001), Longino (2002), and Castel & Sismondo (2003).

An alternative line of argument against the modest principle may be that epistemic oughts are role oughts, which are “oughts that result from one’s playing a certain role or having a certain position” (Conee & Feldman 2004, 175). On this view, they define *standards of good performance*. Conee & Feldman argue that role oughts do not imply can. For example, parents ought to take care of their children, regardless of there being parents who cannot do so. Similarly, in their view, epistemic oughts define people’s standards of good performance *qua* epistemic agents, whether they can meet them or not (2004, 175).

Conee & Feldman do not distinguish, however, between the qualifications required for assuming a role, and the capabilities required for performing it. Not everybody is qualified to assume any role, but those who are qualified and assume a role are expected to be able to perform it. For example, to be a doctor, a person needs to have certain education and pass certain tests, which not everybody can pass. But a doctor’s role responsibilities are defined such that they can be carried out. Similarly, a soldier in a guard post ought to stay alert during her shift, but the shift cannot last 48 hours, since people cannot stay awake for this long. On occasion, an individual soldier may be too tired to perform her duty, but not as a general rule. Indeed, in our society, parenting is regulated such that an adult is normally assumed to be able to perform the role unless proven otherwise, but there are cases in which society, through the legal system, denies people who are found to be unfit parents of their parenting role. The standards by which society judges whether people are competent to assume a parenting role are standards that normal people are generally able to meet.⁹

To conclude, the view that standards of justification required for knowledge do not and should not exceed standards of justified scientific theory acceptance should be generally accepted. The objections to the rationale that underpins it, namely, the epistemic “ought-implies-can” principle, bring down, at most, its extreme version, and not its modest version, which is sufficient for my argument.

6 Scientific Theories – Acceptance versus Belief

So far I argued that well-established scientific theories pass the justificatory standards required for knowledge, although they reflect the social value judgments that were part of the process of inquiry that produced them. One may still deny that they are candidates for knowledge, however, not because they lack justification, but because acceptance, rather than belief is the appropriate cognitive attitude that should be taken toward them. According to this objection, which amounts to denying premise (7) of my argument, since knowing that *p* requires believing that *p* rather than merely accepting that *p*, they are not knowledge.¹⁰

To understand this objection, let us examine the differences between acceptance and belief. Accepting a claim is taking it for granted in one’s reasoning, and it is possible to accept a claim without believing it. Acceptance often results from a consideration of one’s goals, which may be epistemic or non-epistemic, while beliefs are not typically deliberately acquired to advance goals. Acceptance is voluntary, whereas belief is

⁹ For a detailed analysis of the relations between role oughts, epistemic responsibility, and standards of justified belief, see Miller & Record (2013).

¹⁰ The characterization of knowledge as a species of belief is part of the mainstream analysis of knowledge in analytic epistemology. It is not clear, however, that philosophers of science are committed to it, as they tend to focus on acceptance of theories or public claims to knowledge in their normative philosophical analysis, and regard belief as belonging to the realm of the psychology of the scientist, which is not of philosophical interest. For an explicit argument to the effect that knowledge is justified true belief *or acceptance*, see Cohen (1992, 90-92).

involuntary. Belief results in a feeling that something is true; acceptance involves no such feeling (Wray 2001, 325).

In light of these differences, one may object to my reliance on Douglas' argument from inductive risk on the ground that Douglas' argument is about theory acceptance, which is a voluntary decision made on pragmatic grounds, while knowledge involves involuntary belief formation made on epistemic grounds. Hence, it is irrelevant to knowledge (McMullin 1983, 8).¹¹ Thus, it does not follow from the claim that accepted scientific theories inevitably reflect the value judgments and trade-offs that were made prior to their acceptance that beliefs that have the status of knowledge reflect similar value judgments, or so this objection goes.

Let me first address the issue of the alleged voluntary nature of acceptance versus the involuntary nature of belief. It is contested that acceptance is voluntary. According to van Fraassen's influential account of acceptance, for example, accepting a theory is believing that it is empirically adequate (1980, 88). Such second-order beliefs are formed like any other belief, i.e., involuntarily, and may amount to knowledge. But in any case, the objection from the involuntary nature of belief is a red herring, because PE is *not* a theory about when it is appropriate for a subject to form a belief. In particular, PE does *not* state that it is permissible for *S* to form a belief only when it passes the threshold of knowledge-level justification. Rather, PE is a thesis about the factors that determine whether a belief passes this threshold or not. It is permissible in many cases for *S* to believe that *p* even if *S*'s belief that *p* falls short of knowledge. In such cases, *S* will simply have justified belief that falls short of knowledge.

The more serious part of this objection is the claim that value-laden epistemic judgments or attitudes are pragmatic to begin with, and are distinct from value-free epistemic judgments and attitudes that are involved in certifying and regarding hypotheses as knowledge. Lacey (1999, 12-18) provides the most developed suggestion in this spirit. Lacey distinguishes two possible cognitive attitudes scientists may take toward theories: endorsement, and acceptance. Lacey regards endorsement as the right cognitive attitude to take toward provisional less-than-certain theories, while acceptance, or sound acceptance, is the right attitude to take toward certain or near-certain theories. Lacey argues that only endorsement may legitimately involve value-judgments about inductive risks, while sound acceptance should be made on purely evidential grounds.¹² If we follow Lacey's suggestion, only soundly accepted theories are candidates for knowledge, hence scientific knowledge remains value-free.

My reply to this objection is threefold. First, expecting practicing scientists to make such fine-grained distinctions between the different cognitive attitudes they might take toward hypotheses is unrealistic. Lacey's complex framework may be useful to philosophers of science who wish to conceptually analyze and normatively evaluate a case study. But practicing scientists are not philosophers, and cannot be expected to employ it. If we add belief to acceptance and endorsement, we have three possible cognitive attitudes scientists may take toward a hypothesis. If we remember that for each hypothesis, scientists need to determine both their individual attitude and the collective attitude of the community, then for each hypothesis they need to be able to

¹¹ Fantl & McGrath (2010) also interpret Rudner's (1953) argument as concerning pragmatic theory acceptance, rather than knowledge.

¹² Lacey's distinction between endorsement and acceptance is similar to Betz's distinction, which was discussed in §2, between certain or near-certain theories that can be taken to be true for all practical purposes, and less-than-certain hypotheses that should be reported by scientists along with their error probabilities. It fails for similar reasons.

distinguish and correctly choose among six or five possible cognitive attitudes (depending on whether we regard collective belief and collective acceptance as two separate attitudes, or the same one).¹³ Incorporating this fine-grained conceptual framework into actual scientific practice is unrealistic. Moreover, even if scientists learn to distinguish between the different cognitive attitudes, the psychological phenomenon of motivated reasoning sheds serious doubts on their ability to correctly evaluate the level of uncertainty associated with each hypothesis and take the appropriate attitude in each case. Lacey's system violates the modest epistemic "ought-implies-can" principle, which I defended in the last section. Even if we employ a less complex system, expecting scientists to master and use it still seems unrealistic.

Second, the objection according to which value-laden scientific hypotheses should only be pragmatically accepted but not believed rests on a false strict separation between the context of justification and the context of application. It falsely presupposes that values affect only the last stage of research in which hypotheses are accepted or rejected. As we have seen in §4, this picture was debunked by Douglas and others; value judgments and trade-offs between different inductive risks are also made in key stages of inquiry prior to the final stage of theory acceptance, such as data characterization, choice of methodology, data manipulation, evidence assessment, and evidence interpretation. Such stages are common to all scientific research as well as inquiry in general, which often results in beliefs that may amount to *knowledge*.

Complex inquiry in the real world – as opposed to simplified philosophical toy examples – requires making *non-trivial* value judgments. If such inquiry results in beliefs, they will manifest those non-trivial value judgments. Often, a legitimate aim of such inquiry is not only helping make a pragmatic decision, but also yielding knowledge. For example, a police detective conducting a murder investigation may need to rely on forensic techniques to determine whether a suspect's fingerprints or footprints match those that were found in the crime scene. Different forensic techniques have different rates of false positives and false negatives, and the choice between them is therefore inescapably a value choice of a particular balance of inductive risks rather than another. Note that this is a choice of method, which comes at a relatively early stage of inquiry. A legitimate aim of such investigation is not only making the pragmatic decision of whom to charge, but also yielding knowledge. That is, after the detective has properly conducted the investigation and reached the truth, she may legitimately claim to *know* who committed the murder. But this knowledge cannot be separated from the value judgments that were involved in the process of inquiry that yielded it.

Or consider the following example. An editor of a journal in a purely theoretical branch of mathematics sends a paper to two reviewers. They review it and recommend its publication. The editor accepts their recommendation. There is still a chance of error in the paper, which can be reduced if the paper is sent to more reviewers. How is the number of reviewers determined? The editor weighs the risk of publishing an error against other values and practical considerations, such as the speed of publication, availability of reviewers, and overhead costs of additional reviews. It has been empirically found that published papers in mathematical journals occasionally contain errors (Geist et al. 2010; Grcar 2010). Some mathematicians even complain that the rate of published errors is too high (Nathanson 2008). Suppose some editors decided to give higher priority to catching errors before their publication, and raised the reviewing standards, then the published body of mathematical proofs would change. Because new

¹³ For the view that collective belief and collective acceptance are the same cognitive attitude, see Wray (2001); for the opposite view, see Gilbert (2002); for further discussion see Miller (2013).

published proofs draw on old published proofs, the editors' new weighting of values and inductive risks will affect the future proofs published in mathematical journals. *The corpus of accepted mathematical proofs is sensitive to the value judgments of the mathematical community. Some theorems would be considered proven under a certain weighting of values and unproven under another weighting.* Yet it would be absurd to claim that because published proofs in peer-reviewed mathematical journals are not absolutely reliable, or because their reliability may be further improved, they are only candidates for pragmatic acceptance, but not for belief or knowledge. Rather, the knowledge that is obtained by believing such published proofs reflects the value judgments that were made in the process of their acceptance.

Third, hypothesis acceptance in science is not always a purely or even mainly pragmatic matter. Of course, there are cases where theory choice is merely a pragmatic decision, but not all cases are like that. A legitimate aim of science is the production of knowledge for its own sake. For example, the identification of distant astronomical objects such as stars, supernovas, and galaxies is done by computational processing of astronomical data obtained from powerful telescopes. This is basic science as usual. Such computational processing has its levels of accuracy and precision; and its reliability is de facto determined by values such as ease of implementation, cost, and running time, as well as the available hardware (Jarvis & Tyson 1981). Usually, nothing significant hinges on such routine identification of astronomical objects. It is usually done for the pure sake of gaining knowledge. But there might be a case where such identification becomes important from a pragmatic perspective; for example, crucial to successfully carrying out an expensive space mission. In such a case, the computational processing may be amended to reflect the new values and interests – for example, it may be made more accurate at the expense of a slower running time. *The change in values may change the final classifications of some astronomical objects.* But it would be again absurd to deny that the current identification gives us knowledge just because it may change in light of other interests and stakes. That is, it would be absurd to deny that an astronomer may legitimately come to know based on current identifications that a particular astronomical object is, for instance, a distant star rather than a supernova, or to claim that she may only rely on these identifications for making pragmatic decisions but not for forming beliefs and acquiring knowledge.

To sum up this section, granting the objection from pragmatic acceptance amounts to denying that accepted scientific theories and other products of complex inquiry in the real world give us knowledge in many cases they do. In many fields of knowledge, there is no alternative or better way to acquire knowledge than believing our best accepted scientific theories; and we indeed acquire knowledge this way. But such theories already manifest the various value judgments that were made in the various stages of research leading to their acceptance. If we leave out our well established scientific theories from the domain of knowledge and belief, and insist that they belong only to the realm of pragmatic acceptance, the body of knowledge that remains is narrow and trivial. Denying from them the status of knowledge is a price that no reasonable theory of knowledge should be willing to pay.

7 Conclusion

I have argued that knowledge has an inherent pragmatic dimension, drawing on the limitations of inquiry and rational reasoning. Put differently, human beings are imperfect beings producing imperfect knowledge. Their interests and the values they cherish allow them to decide on a non-arbitrary basis which imperfections in their representation of the world they are more willing to bear. The knowledge they end up producing manifests these value choices.

Acknowledgements I thank Arnon Keren, Jacob Stegenga, Dan Hicks, Isaac Record, and Sandy Goldberg for helpful comments on earlier versions of this paper. I thank David Enoch, Moti Mizrahi, and Levi Spectre for useful discussion. I thank two anonymous reviewers for their helpful comments. This paper was presented at the Philosophy of Science (PSA) Meeting (San Diego, 2012), and the Society for Philosophy of Science in Practice (SPSP) Conference (Exeter, 2011). I thank the audience members for useful comments and discussion. This paper was partly written when I was an Azrieli Postdoctoral Fellow at the Department of Philosophy, University of Haifa. I am grateful to the Azrieli Foundation for an award of an Azrieli Fellowship. I am also grateful to the Edmond J. Safra Center for Ethics, Tel Aviv University, and the Dan David Foundation, Tel Aviv University, for postdoctoral fellowships.

References

- Betz, G. 2013. In defence of the value-free ideal. *European Journal for the Philosophy of Science*, 3(2), 207–220.
- Brown, J. 2008. Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167-189.
- Castel, B., & Sismondo, S. 2003. *The art of science*. Peterborough, ON: Broadview Press.
- Cohen, L. J. 1992. *An essay on belief and acceptance*. Oxford: Clarendon Press.
- Cohen, S. 1999. Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57-89.
- Conee, E., & Feldman, R. 2004. *Evidentialism: Essays in epistemology*. Oxford: Oxford University Press.
- Derose, K. 1995. Solving the skeptical problem. *The Philosophical Review*, 104(1), 1-52.
- Douglas, H. 2000. Inductive risk and values in science. *Philosophy of Science*, 67(4), 559-579.
- Douglas, H. 2009. *Science, policy, and the value-free ideal*. Pittsburgh: University of Pittsburgh Press.
- Elliot, K. C. 2011. *Is a little pollution good for you? Incorporating societal values in environmental research*. New York: Oxford University Press.
- Fantl, J., & McGrath, M. 2009. *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Fantl, J. & McGrath. 2010. Pragmatic encroachment. In S. Bernecker, & D. Pritchard (Eds.), *The Routledge companion to epistemology*. London: Routledge (pp. 558-578).
- Feldman, R. 2003. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Geist, C., Löwe, B., & Van Kerkhove, B. 2010. Peer review and knowledge by testimony in mathematics. In B. Löwe, & T. Müller (Eds.), *Philosophy of mathematics: Sociological aspects and mathematical practice*. London: College Publications (pp. 155–178).
- Gettier, E. 1963. Is justified true belief knowledge? *Analysis* 23: 121-123.
- Gilbert, M. 2002. Belief and acceptance as features of groups. *Protosociology* 16: 35-69.

- Goldman, A. I. 1999. *Knowledge in a social world*. New York: Oxford University Press.
- Grcar, J. F. 2010. Errors and corrections in mathematics literature. *Notices of the American Mathematical Society* 60(4): 418-425.
- Hawthorne, J. 2003. *Knowledge and lotteries*. Oxford: Oxford University Press.
- Heil, J. 1983. Believing what one ought. *Journal of Philosophy*, 80(11), 752-765.
- Hempel, C. 1965. Science and human values. In *Aspects of scientific explanation: And other essays in the philosophy of science*. New York: Free Press (pp. 81-96).
- Intemann, K., & De Melo-Martín, I. 2010. Social values and scientific evidence: The case of the HPV vaccines. *Biology and Philosophy*, 25(2), 203-213.
- Jarvis J. F. and Tyson, J. A. 1981. FOCAS: Faint object classification and analysis system. *The Astronomical Journal*, 86(3), 476-495.
- Jeffrey, R., C. 1956. Valuation and acceptance of scientific hypotheses. *Philosophy of Science*, 23(3), 237-246.
- Kitcher, P. 2011. *Science in a democratic society*. Amherst, NY: Prometheus Books.
- Klayman, J. D. 1995. Varieties of confirmation bias. *Psychology of Learning and Motivation*, 32, 385-418.
- Kunda, Z. 1990. The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480-498.
- Lacey, H. 1999. *Is Science value free? Values and scientific understanding*. London: Routledge.
- Lilienfeld, S., O., Ammirati, R., & Landfield, K. 2009. Giving debiasing away: Can psychological research on correcting cognitive errors promote human welfare? *Perspectives on psychological science*, 4(4), 390-398.
- Longino, H. 2002. *The fate of knowledge*. Princeton: Princeton University Press.
- May, J, Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. 2010. Practical interests, relevant alternatives, and knowledge attributions: An empirical study. *Review of Philosophy and Psychology*, 1, 265-273.
- McMullin, E. 1983. Values in science. In P. Asquith, & T. Nickles (Eds.), *PSA 1982, vol. 2*. East Lansing: PSA (pp. 3-28).
- Miller, B., & Record, I. 2013. Justified belief in a digital age: On the Epistemic implications of secret Internet technologies. *Episteme*, 10(2), 101-118.
- Miller, B. 2013. When is consensus knowledge based? Distinguishing shared knowledge from mere agreement. *Synthese*, 190(7), 1293-1316.
- Mizrahi, M. 2012. Does “ought” imply “can” from an epistemic point of view? *Philosophia*, 40, 829-840.
- Nathanson, M. B. 2008. Desperately seeking mathematical truth. *Notices of the American Mathematical Society*, 55(7), 773.

- Neta, R. 2007. Anti-intellectualism and the knowledge-action principle. *Philosophy and Phenomenological Research*, 75(1), 180-187.
- Nickerson, R. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175-220.
- Popper, K. R. 1981. The rationality of scientific revolutions. In I. Hacking (Ed.), *Scientific revolutions*. New York: Oxford University Press (pp. 80-106).
- Pritchard, D. 2006. Review of *Knowledge and practical interests*, by Jason Stanley. *Notre Dame Philosophical Reviews*, <http://ndpr.nd.edu/news/25054/?id=6885>
- Reach, K. 1946. The foundations of our knowledge. *Synthese*, 5(1/2), 83-86.
- Rudner, R. 1953. The scientist *qua* a scientist makes value judgments. *Philosophy of Science*, 20(1), 1-6.
- Schlick, M. 1917/1985. *General theory of knowledge*. Peru, IL: Open Court.
- Solomon, M. 2001. *Social empiricism*. Cambridge, MA: MIT Press.
- Stanley, J. 2005. *Knowledge and practical interests*. Oxford: Oxford University Press.
- Stanley, J. 2007. Replies to Gilbert Harman, Ram Neta, and Stephen Schiffer. *Philosophy and Phenomenological Research*, 75(1), 196-210.
- Stocker, M. 1971. "Ought" and "can". *Australasian Journal of Philosophy*, 49(3), 303-316.
- Van Fraassen, B. 1980. *The scientific image*. Oxford: Oxford University Press.
- Wilholt, T. 2009. Bias and values in scientific research. *Studies in History and Philosophy of Science*, 40(1), 92-101.
- Wray, K. B. 2001. Collective belief and acceptance. *Synthese*, 129(3), 319-333.